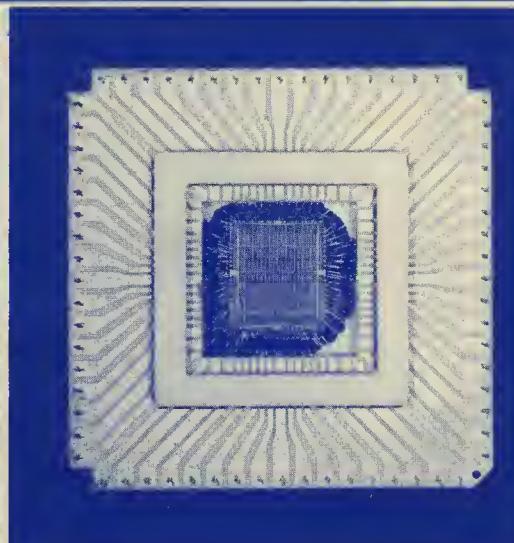
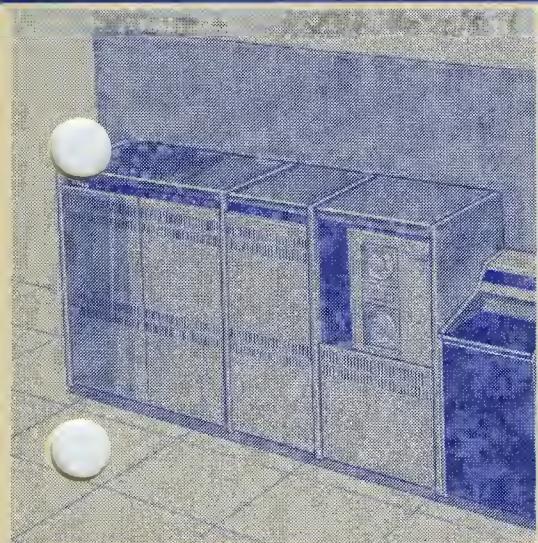


VAX-11/VENUS

VENUS
SYSTEM
DEVELOPMENT PLAN
30 January 1981
Revision 3



COMPANY CONFIDENTIAL

digital

C O M P A N Y C O N F I D E N T I A L

VENUS

SYSTEM
DEVELOPMENT PLAN

30 January 1981
Revision 3

C O M P A N Y C O N F I D E N T I A L

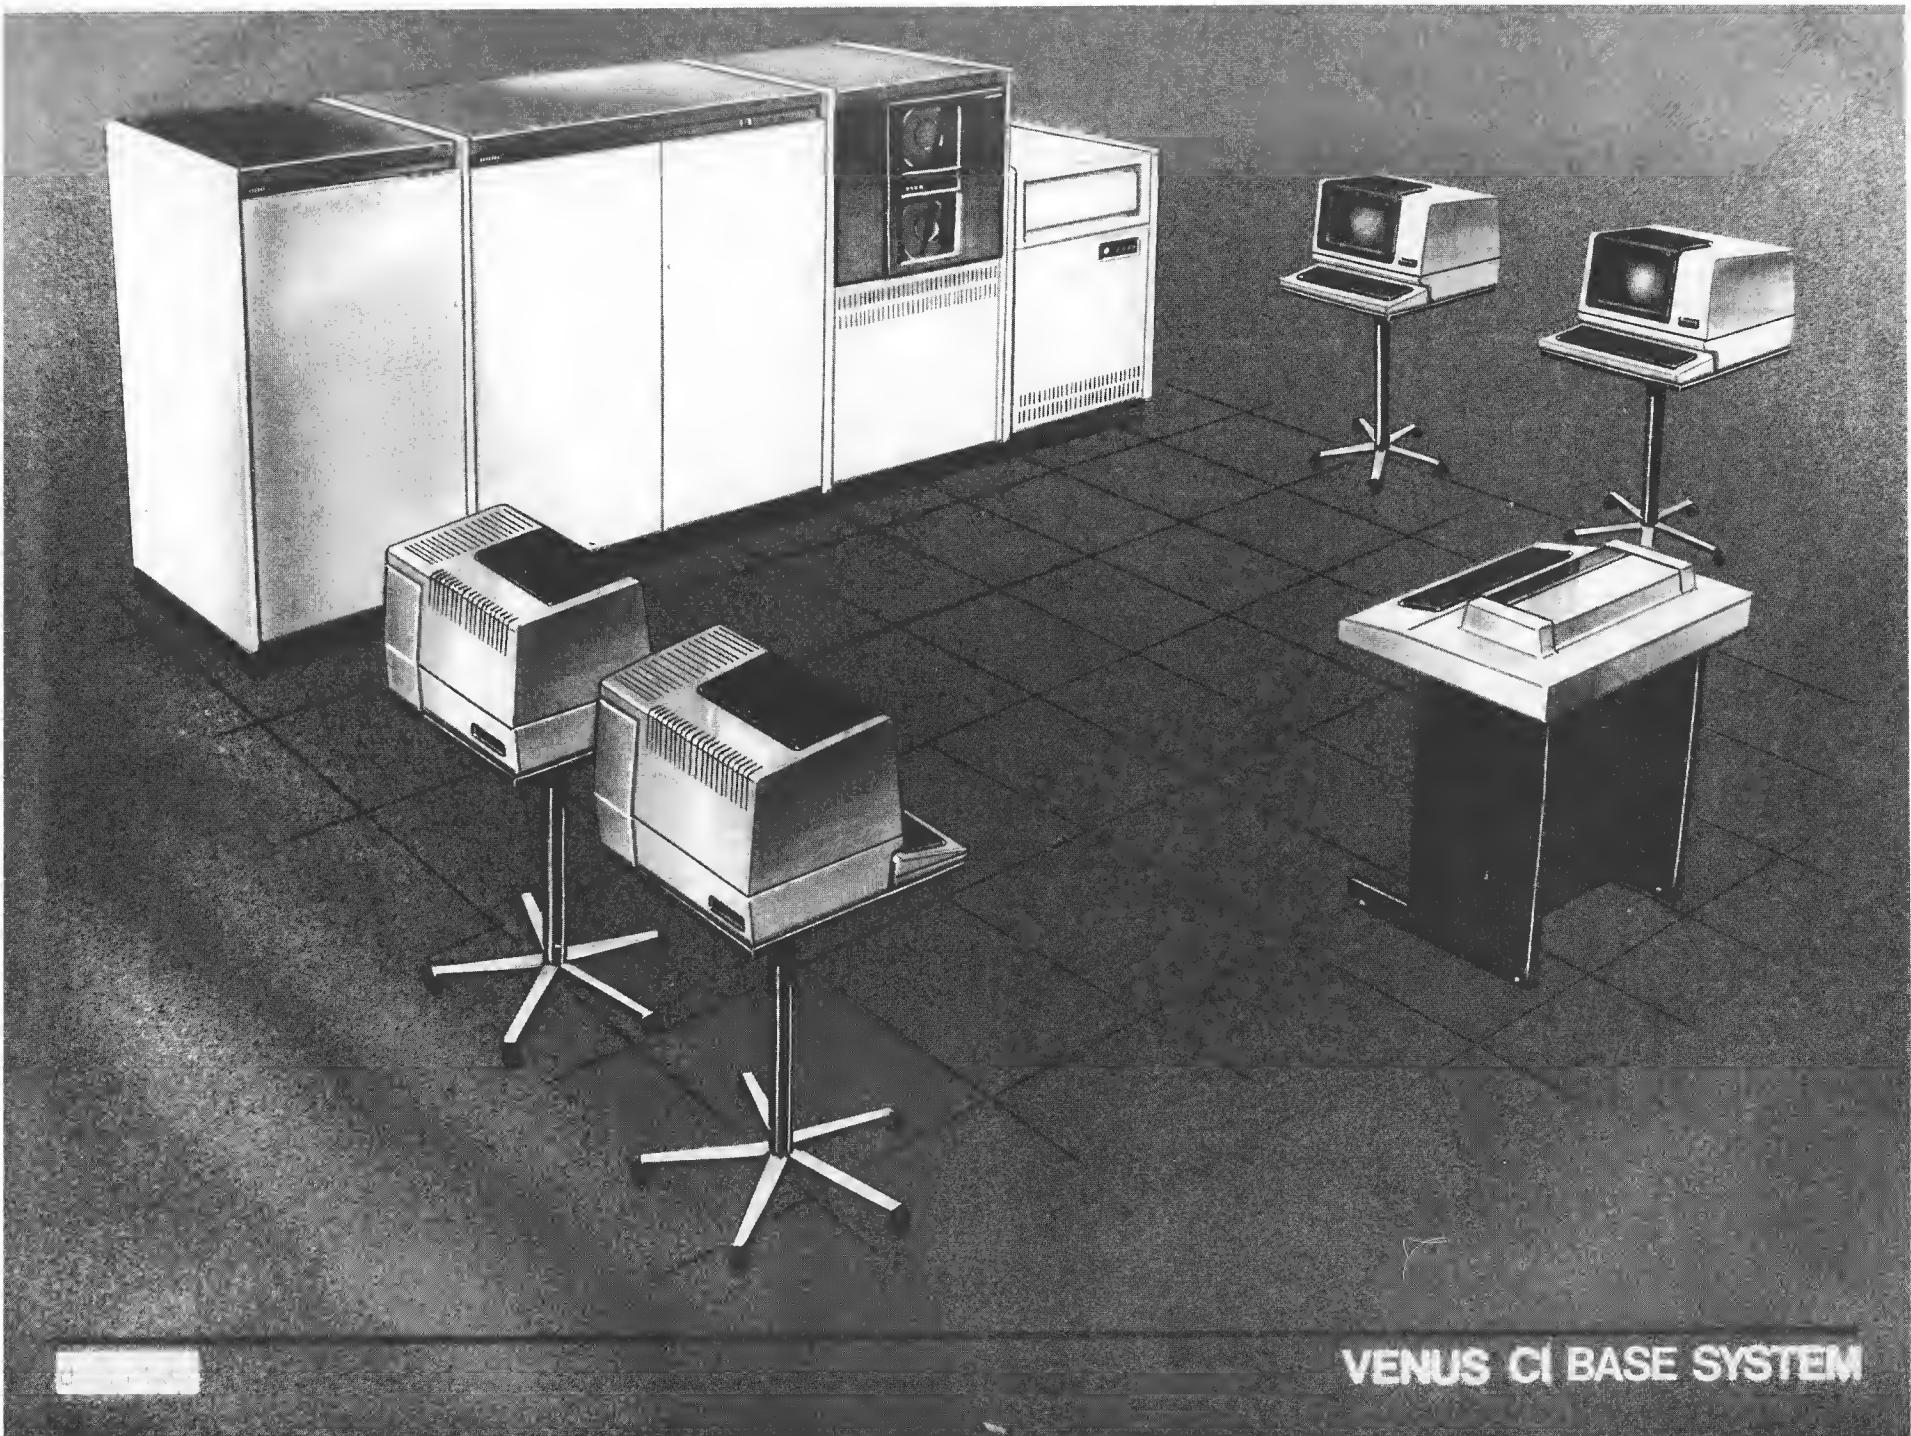
This is the Phase 1 System Development Plan for the Venus Program. It has been reviewed and approved by the development managers, the product managers, and Product Line representatives, and was given formal approval at the Venus Phase 1 Review on December 17, 1980. Venus Product Management's Phase 1 Business Plan is expected to be available in February, 1981, and the Product Contract will be completed in March, 1981.

C O M P A N Y C O N F I D E N T I A L

This document contains confidential information on new products that should be disclosed only to those people engaged in the Program. Under no circumstances should any non-DEC persons be informed about any aspects of the Program or its existence.

Issued by: Vic Ku MR1-2/E47
Editor: Bill English MR1-2/E47

30 January 1981
Revision 3



VENUS CI BASE SYSTEM

VENUS SYSTEM DEVELOPMENT PLAN

This plan presents an exposition of the intended course of the Venus Program. The first section is an executive summary that touches on the essentials of the plan. The next four sections explain what the system is, how it will be developed, how much it will cost, and how the objective of the development program - a marketable system - will be realized. The final section gives more detailed information about the functionality of the various parts of the system.

All of the information given in this plan is necessarily summary in nature, and is derived from the much more detailed plans and specifications for the individual parts of the Program. These detailed documents are listed in Appendix A. Appendix B itemizes the capital equipment to be included in the Engineering breadboards and prototypes. Appendix C lists the product requirements as defined by the Product Lines in conjunction with Product Management and also gives Engineering's response to these requirements. Appendix D identifies all of the people associated with the Program. At the end is a glossary of the myriad mysterious combinations of letters that are tossed about with gay abandon.

The fundamental objective of the Venus Program is to bring a competitive system to market as soon as possible. With this in mind, we are developing an initial system based on the traditional Unibus for communications and unit record, and the new CI-HSC for mass storage (with use of the Massbus for a backup system). The CI project and its peripherals are well-enough along so we can be confident that a system based on it can be delivered in the timeframe set for the Program. The NI however is not yet at a sufficient confidence level - many projects are still only being planned. On the other hand, as soon as the initial system is wrapped up, the Venus design team will go to work on a mid-life kicker, involving not only conversion to NI but implementation of advanced packaging as well.

This plan has been approved by the Venus managers and supervisors whose signatures appear on the next page.

By signing this plan each of us indicates, in our best judgment, that

- 1) I understand this plan and feel that I fully appreciate its implications both for myself and for the Venus Program;
- 2) I am aware of the expectations the Venus Program has of my group, and I am confident we can fulfill those expectations;
- 3) This plan fully addresses all expectations I have of the Venus Program, and I feel confident the Program can fulfill those expectations;
- 4) I am confident that, overall, the objectives of this plan are achievable in the timeframe indicated; and
- 5) I will communicate, in writing, to the Program Manager any specific reservations that I have about any item in the plan.

Engineering Group Manager
George Hoff

George Hoff
.....

Program Manager
Vic Ku

Vic Ku 24.7.180
.....

Product Manager
Carl Gibson

Carl Gibson
.....

System Architect
Jud Leonard

Jud Leonard
.....

CPU and Adapter Engineering
Sas Durvasula

Sas Durvasula
.....

Peripheral Engineering
John Bloem

John Bloem
.....

Technology Engineering
Sultan Zia/Jim McElroy

Sultan Zia/Jim McElroy
.....

Diagnostic Engineering
Dick Beaven/Dale Cook

Dick Beaven/Dale Cook
.....

VAX/VMS Engineering
Joe Carchidi/Chuck Samuelson

Joe Carchidi/Chuck Samuelson
.....

Customer Service
Walter Manter/Mike Robey

Walter Manter/Mike Robey
.....

Manufacturing
Charlie Bradshaw/John Grose

Charlie Bradshaw/John Grose
.....

MCA/LSI Manufacturing
Ken Brabitz/Bill Walton

Ken Brabitz/Bill Walton
.....

Qualification Engineering
Ron Setera

Ron Setera
.....

Educational Services
Ed McFaden

Ed McFaden
.....

Site Engineering/Operations
Roy Rezac/Nick Cappello

CONTENTS

1	THE VENUS PROGRAM	1
1.1	Program Priorities and Highlights	1
1.2	Why Venus?	8
2	THE SYSTEM	12
2.1	Logical Organization	12
2.2	Physical Organization	15
2.3	Peripheral Equipment	20
2.4	Software	23
2.5	Performance	24
2.6	RAMP and Data Integrity Features	28
3	DEVELOPING THE SYSTEM	31
3.1	Development Strategy	31
3.2	Development Tools	37
3.3	Phase Review	39
3.4	Program Tracking and Major Milestones	40
3.5	Risks and Dependencies	45
3.6	Time to Market, System Configuration, and Supported Peripherals	49
4	COST	52
4.1	Development Cost	52
4.2	Manufacturing Cost	55
4.3	Life Cycle Cost	58
5	SYSTEM REALIZATION STRATEGIES	59
5.1	Manufacturing Strategy	59
5.2	Dock Merge	60
5.3	Qualification Plan	62
5.4	Diagnostic Strategy	63
5.5	Customer Service Plan	65
6	FUNCTIONAL CHARACTERISTICS	68
6.1	CPU Cluster	68
6.2	Peripherals	70
6.3	Software	72
6.4	Technology	74
Appendices		
A	System Related Plans and Specifications	80
B	Breadboard/Prototype Capital Equipment	82
C	Product Requirements - Engineering Response	85
D	Who's Who in the Venus Program	93
Glossary		98

1 THE VENUS PROGRAM

This section summarizes the elements of the Program and explains why Venus is the best candidate to carry the Digital logo at the high end of the VAX line. Here and elsewhere in this plan, the discussion of schedules and commitments employs such terms as FRS and FVC; these terms are explained in the Glossary.

1.1 PROGRAM PRIORITIES AND HIGHLIGHTS

At the core of the Program is an engineering development plan organized around a set of priorities or overall goals. Within the framework of these goals, we have defined a product using state-of-the-art technology.

Program Priorities

The development strategy is geared to these goals listed in order of importance to the Program.

Maximize customer satisfaction (superior to comparable IBM systems)

Minimize life cycle cost

Cost of ownership less than comparable 11/780 systems (we believe we are achieving this, but the model is unclear)

Meet transfer cost targets: comparable to equivalent 11/780 systems (\$73K for the CI Base and < \$52K for the IDTC Base defined on the next page vs. \$48K for the VAX11/780SV-AXCVA)

Maximize price/performance gain over 11/780 (4 times 11/780 performance except floating point timing without FPA)

IO architecture based on new Corporate interconnects (CI at FRS, NI on Unibus as soon as possible)

SBI capability for 11/780 migration

FRS by 7/83; total number of systems shipped in first two years: 760 in FY84, 2880 in FY85

Full-featured system to ship no later than nine months after FRS of initial midrange system

Design system for maximum dock mergeability

Maximize RAMP features

Minimize development cost

Naturally these various goals conflict with one another. A significant part of the decision making reflected in this document and that will be exercised throughout the development is determining what tradeoffs to make among these goals to create the most viable product.

System Definition

Listed here are the fundamental constituents of the three categories of system upon which the Program is based. The CI Base defines the principal system configuration and is expected to be the basis for all midrange and larger systems. The IDTC Base, which Marketing may not make available until almost a year after FRS so as not to compete with the 11/780, has a disk-tape controller within the CPU cabinet and is the basis for the smaller systems, although it too is expandable. Note that the integrated disk-tape controller is composed of current products - DW780 and UDA50 - just with new packaging and a special backplane; it involves no new hardware, and no new VMS software is needed. The Massbus Base is a backup at FRS for the CI-based systems, and it will be available six months later for 11/780 upgrade. Note that the DZ730, the "Combo" board, has controllers for eight asynchronous lines, one synchronous line, and one line printer; which controllers are actually used depends on what other equipment is attached and what arrangements are made in terms of distribution panels for the various lines. All three system types in turn contain a set of common elements that we shall first define as the "CPU cluster".

CPU cluster

Processor (including power)
1 LA120 console terminal
1 RL02 console load device
1 megabyte of MOS memory
1 SBI adapter (SBIA)
1 DW780 Unibus adapter
1 DZ730 Combo (8 asynchronous lines)

CI Base

CPU cluster
3 additional megabytes of MOS memory
1 additional DZ730 (16 asynchronous lines total)
1 CI780 CI adapter
1 HSC50 controller
1 RA81 disk
1 TA78 tape

IDTC Base (Integrated Disk-Tape Controller)

CPU cluster
1 DW780-UDA50 integrated disk-tape controller
(4 disk ports and 1 tape port)
1 Pinon or RA81 disk
1 LCGCR tape

Massbus Base

CPU cluster
1 additional megabyte of MOS memory
1 additional DZ730 (16 asynchronous lines total)
2 RH780 Massbus adapters
1 RP07 disk
1 TM78-TU78 tape

From these basic configurations, the Venus Business Plan defines a number of dock-mergeable packaged systems for handling the largest volume markets at the lowest cost. Unless otherwise specified, information in the present document pertains generally to the CI Base and systems containing it. The term "system base" refers to the base of any type of system.

The basic DMT system is the CI Base minus disk and tape. The CPU kernel (for DMT and similar purposes) is the processor plus 1 MB of memory.

Program Highlights

The System

High Availability

CPU cluster - 99.5%

System base (including software) - 98.5%

No more than 1 software crash per month

MTBF (mean time between failures)

In each case the second number is the MTBF perceived by the customer, taking into consideration the fault tolerance of the CPU kernel and scheduling maintenance during off time. Figures include field performance data for the LA120 and RL02, but in Venus these devices will have a lower duty cycle, which will increase the MTBF somewhat.

CPU kernel (DMT) - 1790/1967 hours

CPU cluster - 895/943 hours

CI Base - 411/423 hours

IDTC Base - 597/618 hours

Massbus Base - 463/476 hours

Hydra-type system configuration possible (Venus satisfies the established criteria for use as a node in Hydra, although no qualification plan has yet been developed)

System MTTR - 3 hours; MDT - 5 hours

System installation and acceptance < 48 hours

Warranty cost: CPU cluster - \$5823

CI Base - \$14,399

IDTC Base - \$9519

Massbus Base - \$15,799

Maintenance cost (at 20th quarter of shipments)

CPU cluster - \$401 per month

CI Base - \$1237 per month

IDTC Base - \$689 per month

Massbus Base - \$965 per month

BMC (basic monthly charge) - TBD in Phase 2

Advanced RAMP features

Instruction retry



digital

VENUS CI BASE SYSTEM

Multiplexers built into terminator chips for diagnostic inspection of all backplane signals

Diagnostic resolution to module in at least 95% of solid failures; resolution to chip level for RAM failures

All MCA chips and most RAMs mounted in sockets to allow field replacement

Integrated, intelligent maintenance/operator console (T-11 based)

Loopback diagnostics for IO controllers

User mode diagnostics

Etch backplanes (> 90%)

Early warning on low voltage/high temperature

Power fail recovery

Remote diagnostic link

Battery backup on time-of-year meter (100 hours) and memory (10 minutes)

The Technology

Press pin, multilayer, controlled impedance backplanes (16 layers)

8-layer, controlled impedance L-type modules (4 signal layers)

Macrocell arrays (MCA) - Motorola Mosaic I ECL gate array LSI technology

Digital Hudson plant second source for MCAs

ECL 10K MSI/SSI chips

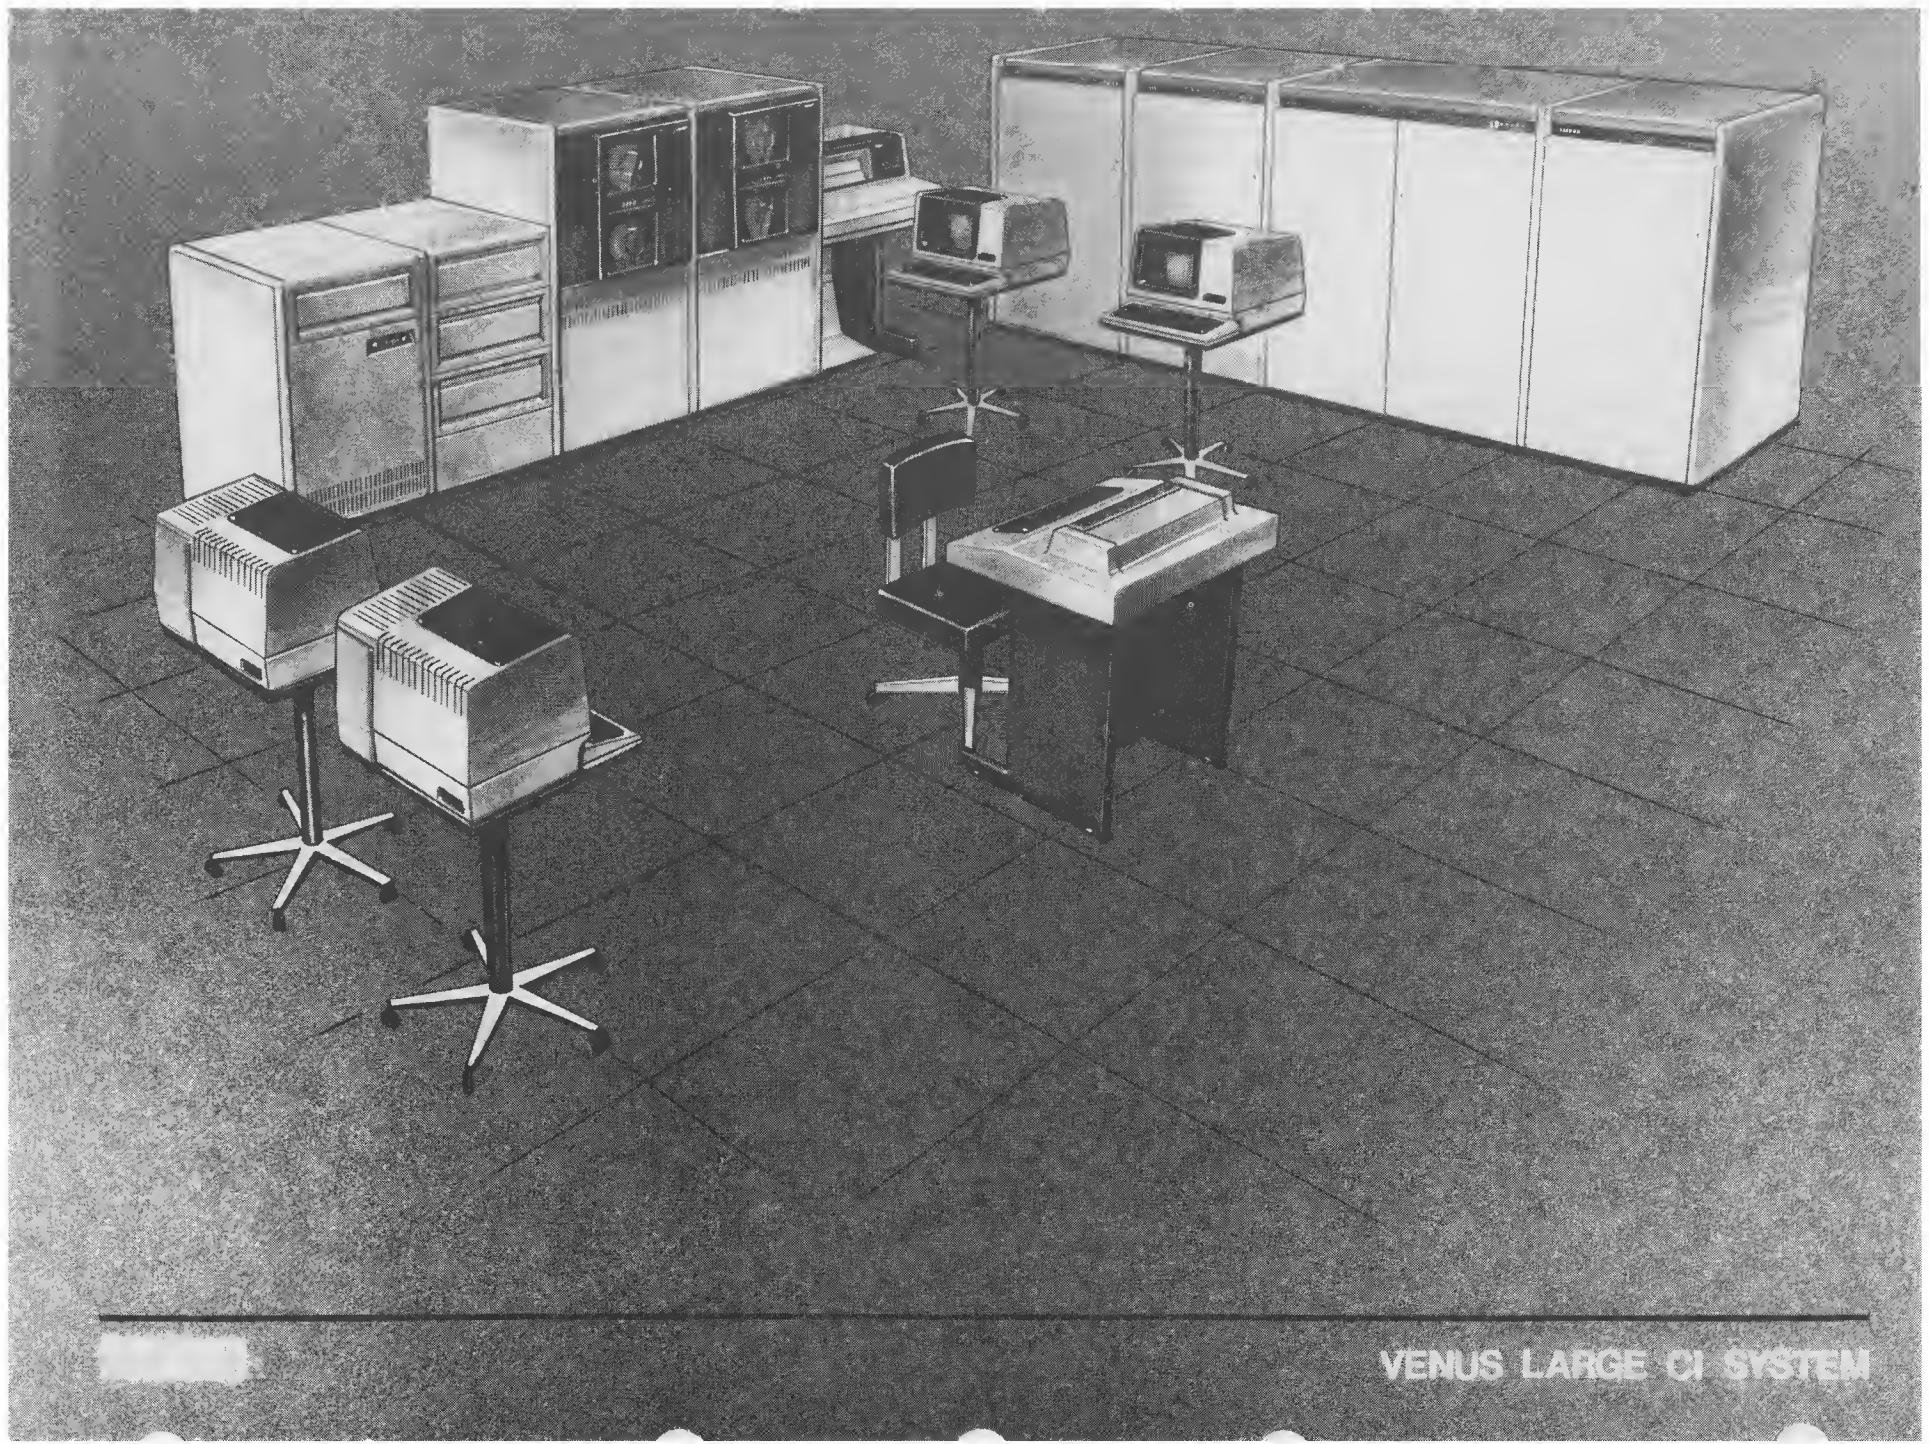
Air cooled

Individual heat sinks mounted on each MCA

1K and 4K ECL RAMs

64K MOS memory chips (16K backup for breadboard)

Corporate "switching regulator" modular power supplies



VENUS LARGE CI SYSTEM

The Product

High end VAX system - corresponds exactly to the VAX architecture defined in the VAX System Reference Manual (DEC Standard 32); the AXE program, VMS and layered products will be used to verify the architecture

Supplements the VAX11/780 RP06-TE16-based system (codes SV-AXCVA-LA and LD) and all larger 11/780 systems; it will be the primary offering for applications requiring greater than 1 MB of memory and 200 MB of disk

Many configurations dock mergeable

Absolute compatibility with VAX/VMS and layered software products for equivalent hardware (at present the medium for software distribution is an open issue, and a task force is investigating the matter vis-a-vis the entire VAX family)

16K byte writeback cache with ECC

4 times VAX 11/780 performance

Instruction byte prefetch

Custom register file for A bus adapters

8K writable control store (there is no WCS option, but 1K is available to the user with limited support comparable to that on 11/780)

Optional floating point accelerator, 4 times 11/780 FPA performance

MOS memory in 1 MB increments to 32 MB (8 at FRS)

Memory data register chip

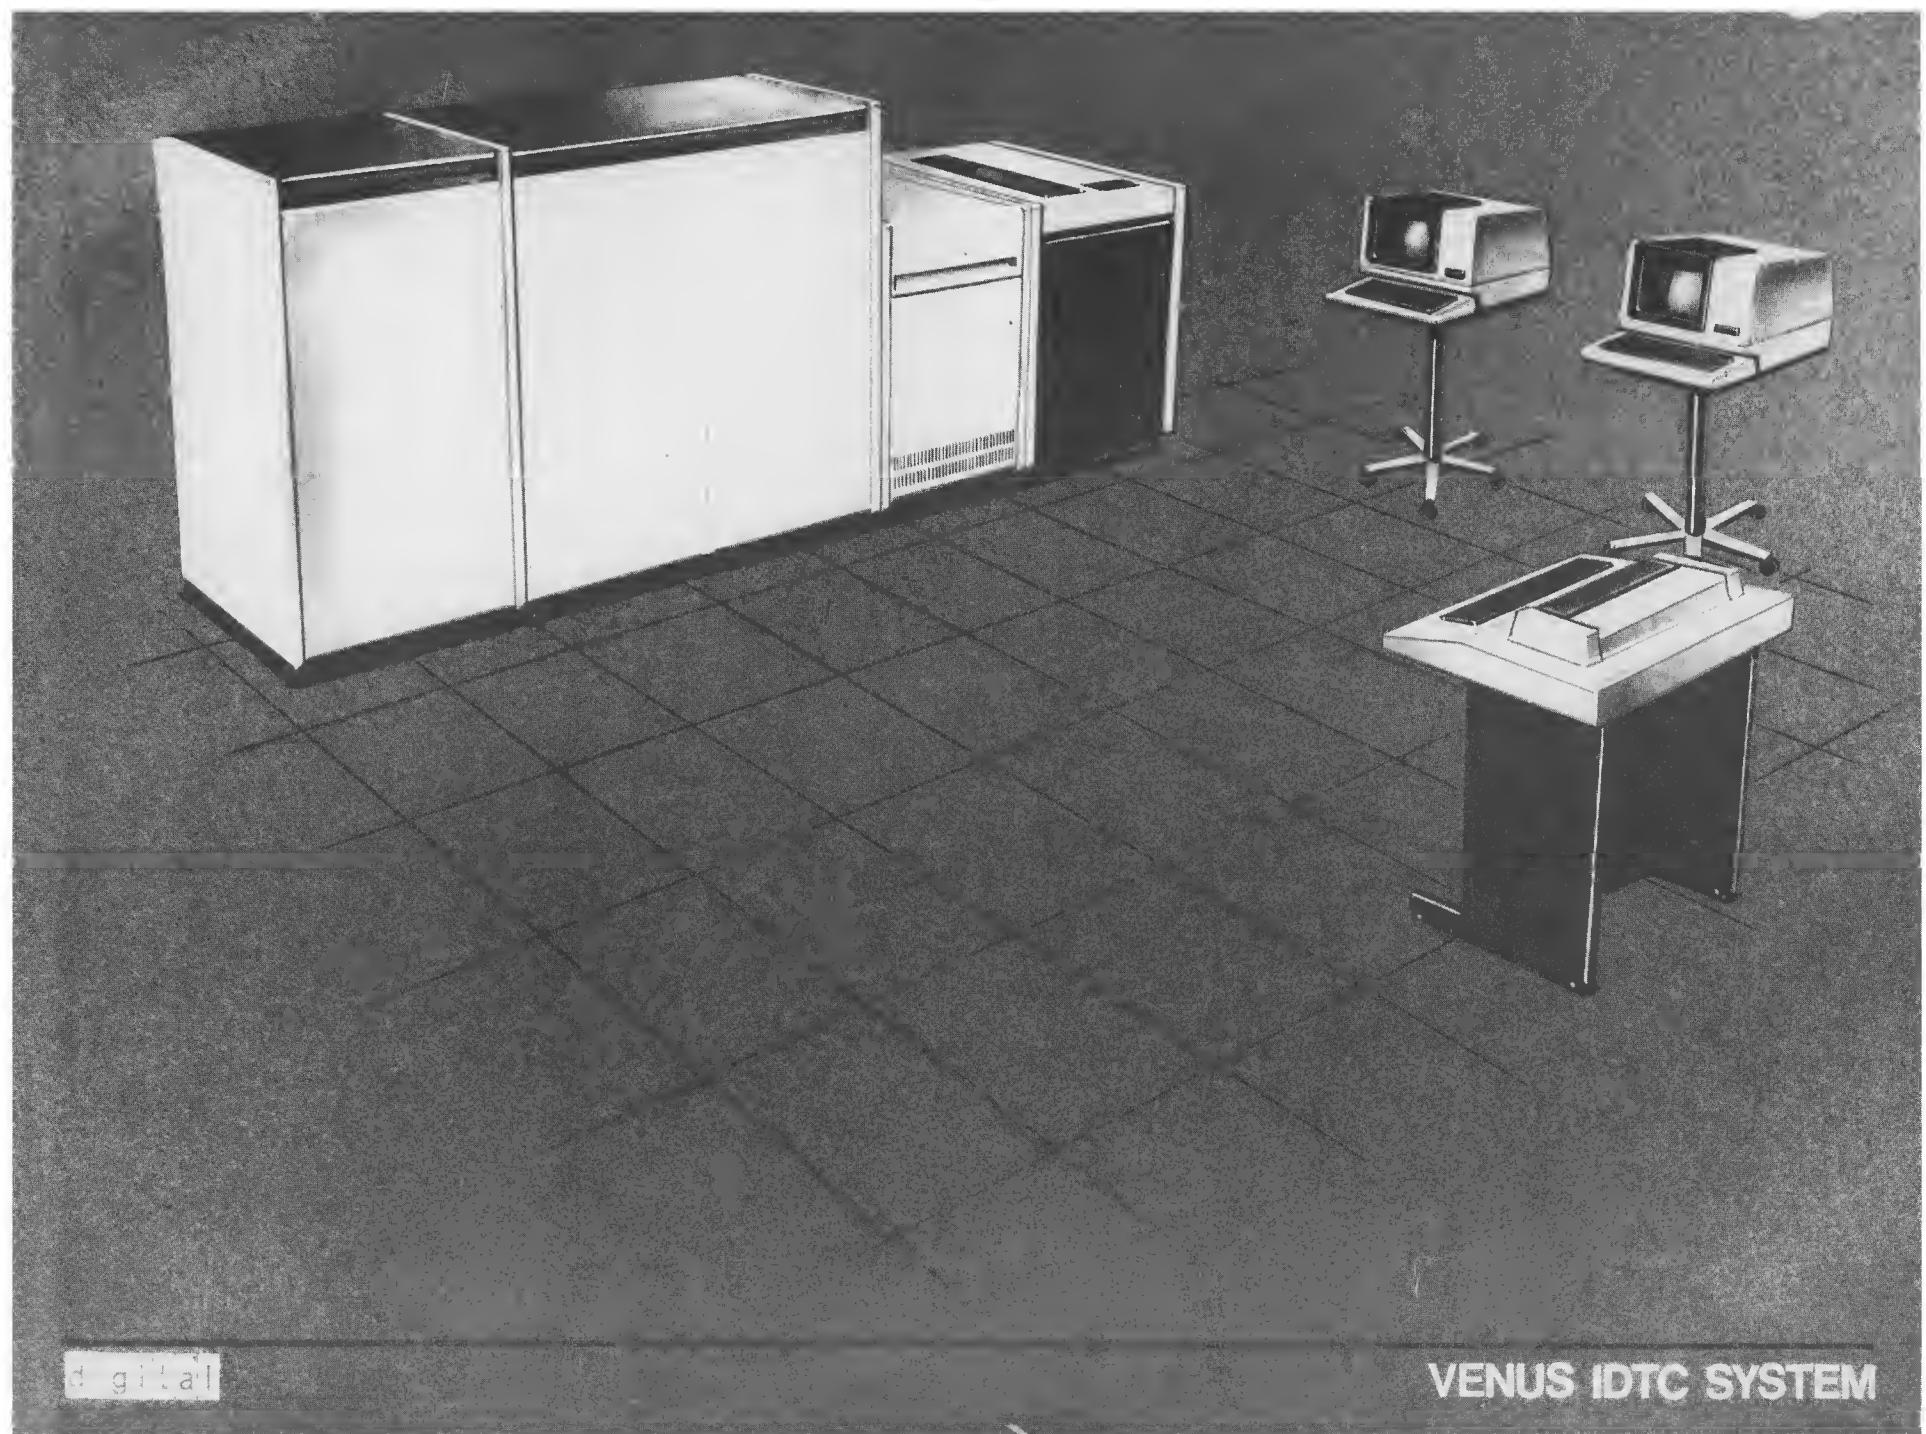
Supports new CI interconnect - in fact the Program depends on the CI for HSC50 and Hydra

3 taps on adapter bus: 2 for SBI adapters and 1 (single slot) for Product Lines or special equipment (second SBI available 3 months after FRS)

VMS support for Venus Processor initialization, error handling, IO adapters and console

Supports PDP-11 compatibility mode

High volume fabrication, assembly and test



digital

VENUS IDTC SYSTEM

Optimized cooling/packaging for class A environment (air conditioned, 15-32 degrees C, relative humidity 20-80%; raised floor prefered but not necessary

Meets or exceeds requirements of DEC Standards 60 and 102, including RFI/EMI, UL, CSA, IEC, VDE (we are also working with the Corporate Central RFI/EMI Group to ensure compliance with new FCC regulations)

Power consumption

CPU cabinet: 5.2 kW, 17,700 Btu/hour, 8.7 kVa (CPU, FPA, 8 MB, console, 1 SBIA, 2 DW780s, CI780)

Unibus-console cabinet: .8 kW, 2700 Btu/hour, 1.3 kVa (RL02, BAll, 2 DZ730, etc)

Memory expansion cabinet with 24 MB: 2.5 kW, 8300 Btu/hour, 4.1 kVa

System fully DMTed and PMTed

Meets or betters European noise standard (60 dbA), and we will deal aggressively with vendors concerning noise requirements for peripherals

Will provide a strong functional base for:

- High end real time
- Foreign device connect
- Transaction processing
- Timesharing
- Batch
- Distributed processing
- Distributed data base management

When and How Much?

The individual Venus managers and supervisors have identified the various tasks and events that make up the overall Program and have determined reasonable schedules for implementing them. Combining these schedules, taking into account the interdependencies among them, results in these dates that guide the Program.

CI Base	FRS July 1983
Transfer cost: \$73K	

Floating Point Accelerator	FRS July 1983
Transfer cost: \$3.5K	

Massbus Base	FRS January 1984
Transfer cost: < \$71K	Backup July 1983
Dual SBI System	FRS October 1983
Transfer cost: configuration dependent	
Memory Expansion Box	FRS January 1984
Transfer cost:	
Up to 12 MB: \$6900 + \$2004/MB	
More than 12 MB: \$7900 + \$2004/MB	
IDTC Base	FRS May 1984
Transfer cost: < \$52K	

The current schedules for the CI780, HSC50 and CI peripherals support the given FRS date for the CI Kernel. However, four to six quarters before FRS, the status of these projects will be reviewed, and a decision will then be made whether to continue on the present course or switch to the backup system. If the latter becomes necessary, we will nonetheless endeavor to bring out the CI system as soon as possible, depending on availability of the needed options.

It should be noted that the above are the times at which the Program will have the various products ready for FRS. Because of Marketing considerations however, Product Management may in some cases prefer actually to make them available at a later time.

1.2 WHY VENUS?

In every category - from market suitability to customer satisfaction, from cost/performance to expansion capability - Venus is the computer for the high end of the VAX line in the mid-Eighties. It is also consistent with the Corporate strategy of settling principally on a 32-bit architecture by 1985.

Leadership

In terms of performance, availability, cost of ownership, and range of applicability, Venus is a major stride both within Digital and in the industry as a whole. These exceptional improvements are due to use of state-of-the-art Mosaic I ECL array technology.

Venus will be the most powerful system in Digital's VAX product offering, with higher availability and lower cost of ownership than any comparable Digital system. With the layered

software products planned for delivery in the early Eighties, Venus will meet the needs of a broad range of technical, commercial, and special-application customers. By incorporating the latest technology in hardware and software, Venus will bring "people oriented" computing to performance levels never attained before.

Cost/Performance

Venus-based systems will exceed the 22% per year cost/performance improvement shown by the computer industry as a whole. Thus in raw cost/performance terms, Venus can be expected to be as competitive in the mid-Eighties as the 11/780 is today; furthermore, in functionality terms, our large software development for VAX systems will make Venus systems even more attractive.

Performance

On computation-limited workloads, Venus will have 4 times the throughput of a comparable 11/780. On IO-limited workloads, the improvement will depend principally on the capability of the HSC50. Within the mechanical constraints (and most IO limits are mechanical), the SBI-CI-HSC50 subsystem is capable of considerable data-transfer optimization; but beyond this, the HSC50 also has features for optimizing the mechanical operations of the disk itself.

Cost

The cost of ownership of Venus systems will be equal to or better than comparably configured 11/780 systems. "Comparably configured" means that Venus main memory and disk capacity are three to four times that of "comparable" 11/780s. We expect to reduce FA&T costs by dock merge of many systems and offering packaged systems, and to minimize life cycle cost by careful design of the system, its RAMP features, and our service and manufacturing strategies. As an example the total CPU cluster warranty and maintenance cost for a representative sample of 13,710 systems will range from \$91.6 to \$102 million depending on the extent to which sockets are utilized. Hence maximum socket use will save \$10.4 million over the life of the product.

Market Suitability

Venus will span two very different market places: as a small mainframe, it will be comparable in performance to a 370/168; as a high end minicomputer, it will pick up the real time, interactive, and distributed processing applications as they grow to higher throughput requirements. Although initially most appropriate to midrange scientific computation markets, Venus will be able to take full advantage of VAX/VMS software efforts to penetrate commercial ADP applications. Over its life, Venus is expected to be installed in roughly as many commercial as scientific applications.

Timeliness

The 11/780 is presently under fire from many directions - the IBM 4331 and 4341 (and to-be-announced low-end H-series system), the Interdata 3240, and expected offerings from DG, SEL and HP are all aggressive products in the 11/780 market space. Venus's market introduction in 1983 will provide Digital with a resource to meet these challenges.

Flexibility

While Venus's high marks in performance, cost and other characteristics make it an excellent instrument for attacking new markets and attracting new customers, its compatibility not only with VAX/VMS and all VAX layered software products, but also with the PDP-11 will make it exceptionally attractive to present customers for upgrading and networking.

Customer Satisfaction

Low component count, design that anticipates exceptional conditions, and extensive checking circuitry give Venus high inherent reliability. Ride-through strategies to survive transient errors help minimize vulnerability to intermittent failures, and error logging gives Customer Service the information to locate and repair such faults. The diagnostic logic supports module-level fault isolation, and in many cases isolation to chip level. These and other RAMP features contribute directly to customer satisfaction, as does the resulting lower life cycle cost.

Interconnectability

Venus will interface to existing and anticipated applicable Corporate interconnect mechanisms. Easy access to traditional Unibus and Massbus devices will be available.

Multicomputing

Venus processors can be loosely coupled via the CI, e.g. the Hydra configuration. With full exploitation of the Corporate interconnect strategies, Venus can lead the real computer revolution of the Eighties - using arrays of \$350K machines to solve multimillion dollar problems.

2 THE SYSTEM

A Venus System is more than just hardware - it is also the software that makes the system go, and its performance characteristics and RAMP features, all of which we consider here.

2.1 LOGICAL ORGANIZATION

In the block diagram of the central part of the system, Figure 2.1, the processor comprises the five blocks interconnected by the diagnostic bus: the single block in the upper left and the upper four blocks in the center column. These four blocks represent the instruction or I box, the execution or E box, the floating point accelerator FPA (also called the F box), and the memory control or M box. The last of these provides the connecting link to both memory and the IO subsystem, whereas the upper three blocks comprise the "processing" part of the processor.

The heart of the entire system is the instruction box, which receives the instruction stream of bytes from memory, and from it determines what other information to retrieve and what activity to initiate in the E box. The execution of each instruction is done in four stages, with the I box handling the first three: fetching the instruction, calculating the required addresses, and fetching the operands. It then turns the execution of the instruction over to the E box, but at the same time it helps to speed up overall operation by starting to work on the next instruction. The E box, based on a binary/BCD ALU, carries out whatever logical, arithmetic and other operations are required to execute the instruction, after which it sends to the I box any results that are to be written in memory. If the optional FPA is included, the E box uses it to speed up the execution of floating point instructions, but from the point of view of the I box, the FPA is simply an extension of the E box. Basic to

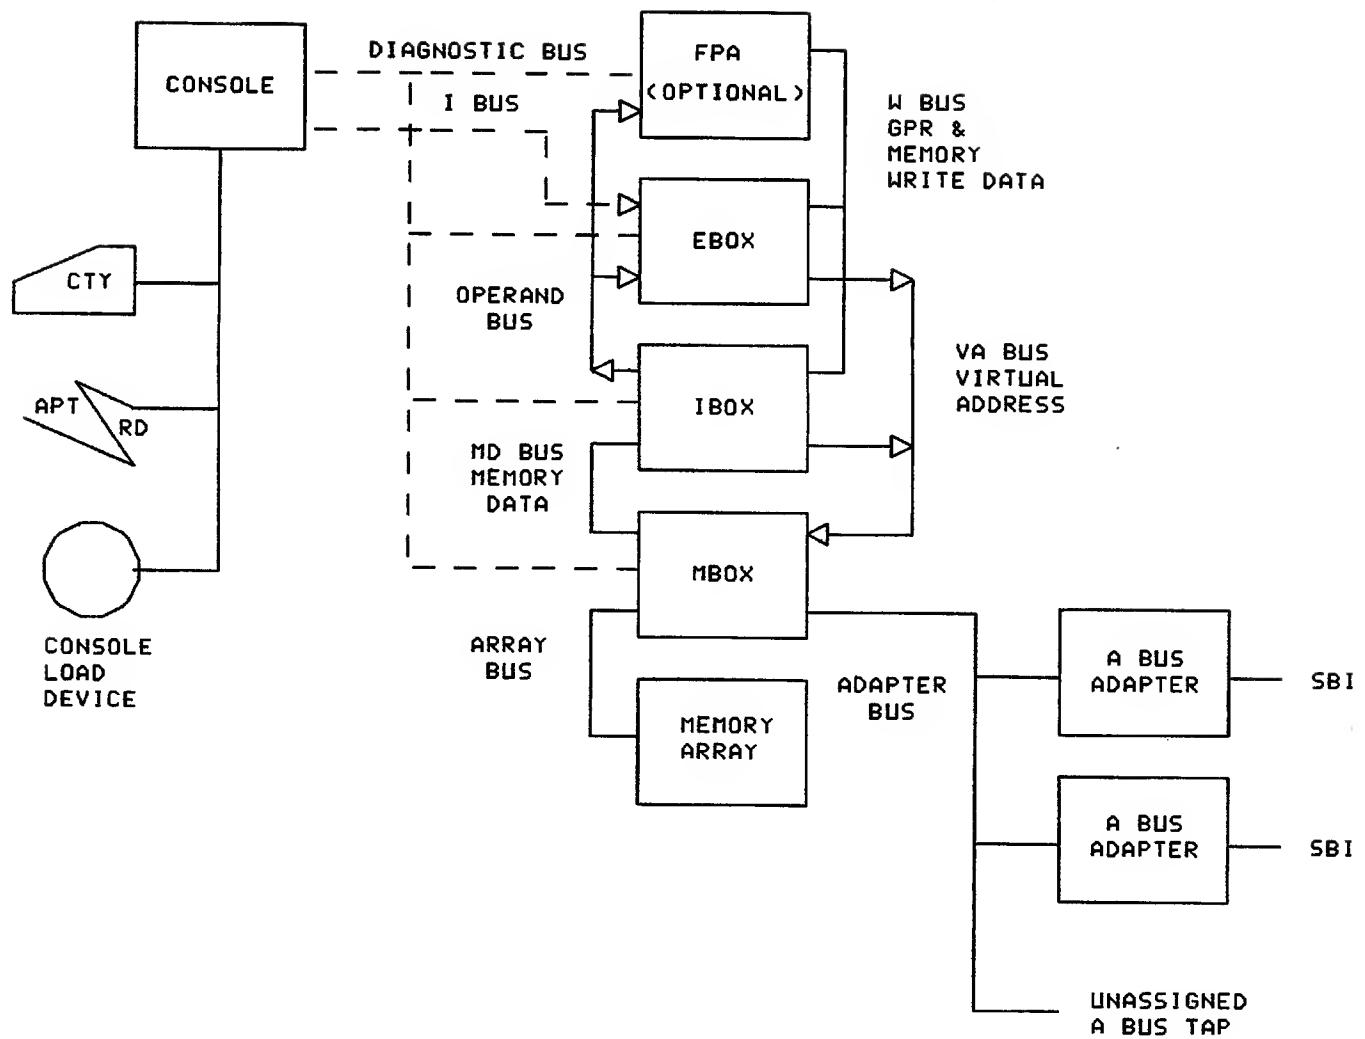


FIGURE 2.1 VENUS PROCESSOR LOGICAL ORGANIZATION

the accurate and fast manipulation of data are the general purpose registers (GPR), of which each of the three boxes has one or two sets of sixteen. Altogether five copies of the GPRs are kept to guarantee very fast and flexible access and instruction retry.

Interconnecting the various boxes are a number of buses. All movement of data between the processor and both the memory array and the IO subsystem occurs via the MD bus that connects the M and I boxes. Over this bus the I box receives the I stream and the memory operands. It passes the latter on to the E and F boxes over the operand bus. Results from either of these boxes are sent via the W bus to the I box, which in turn passes them on to the M box over the MD bus. The W bus is also used for keeping the five sets of GPRs identical to one another. Both I box and E box can supply addresses (almost always virtual) over the VA bus to the M box. All buses and registers handle 32-bit longwords.

Also contained in the processor is a microprocessor-based console, which is connected to all four of the boxes by a serial diagnostic bus. The console provides the system clock, a time-of-year clock, and environmental monitoring. Associated with the console are a local LA120 terminal for use by the operator, an RL02 removable disk (mounted in the Unibus cabinet) for bootstrapping and diagnostic activities, and a remote diagnostic link which can be utilized by an APT window. Bootstrapping is done by the console passing initializing and setup information in bytes to the various boxes over the diagnostic bus. The 8K x 84 bit control store for the microcode is in the E box, but each of the other boxes has a small control store to hold special microcode for its own operations. Also connecting the console to the E box is the C bus for communicating with the software and performing console functions.

The M box includes a 16K-byte data cache, error detection and correction circuits for the cache and memory array, and microcode-driven control logic for governing communication with the IO subsystem as well as handling memory. The control part includes special byte write logic to speed up the insertion of bytes in longwords, and refresh logic for the MOS RAMs in the array. Connection to memory is via the array bus, and to the IO subsystem via the adapter or A bus. Each array board contains one MB of MOS storage; a typical memory is one to eight MB, with expansion to 32 MB possible. It is expected that within two years after FRS, memory density will improve by a factor of four (the addressing capability of the processor hardware is 512 MB). On the A bus are adapters for one or two

SBIs. IO bandwidth is significantly increased over 11/780 by removing all CPU-memory traffic from the SBI. There can be no SBI memory, and there is no support for a device like the MA780.

Although the FPA is optional and memory size is variable, most variation from one system to another occurs in the IO subsystem. Such variation is considerable, including both the interconnects and a large variety of peripheral equipment, but in general system configurations are of two fundamental types, based either on the CI Base or the IDTC Base. The basic constituents of these systems are shown in Figure 2.2. Available initially will be systems built on the CI Base, where the IO subsystem has an SBI with units that interface to a Unibus and a CI. The latter has fifteen nodes that allow connection to HSC50 disk or tape systems, and to other computers in a multicomputer system. For a smaller scale system, the IDTC Base has a DW780-UDA50 disk-tape controller in place of the CI780. Either system can have a second Unibus adapter in the CPU cabinet; more adapters, including the RH780 and DR780, can be added outside the cabinet; and a second SBI can be installed to handle the external adapters and thus share the load. The Massbus Base, which is the backup system, has two DW780s in the CPU cabinet, and requires an expansion cabinet to accommodate the Massbus adapters.

2.2 PHYSICAL ORGANIZATION

Figure 2.3 shows the physical layout of the CPU cabinet. In terms of general organization - position of backplanes, power supplies, cables, blowers and the like - the layout is the same for all systems. Viewed from the front, the left half of the module area is a single CPU card cage containing two backplanes for memory and processor, both implemented in L-type modules. The left backplane accommodates eight memory array boards of one MB each. Beside the memory is the processor, which if the optional floating point accelerator is included, requires seventeen modules.

CPU Box	Number of Modules
---------	-------------------

I box	3
E box	4
System clock	1
Control store	2
M box	3
FPA	3
Console	1

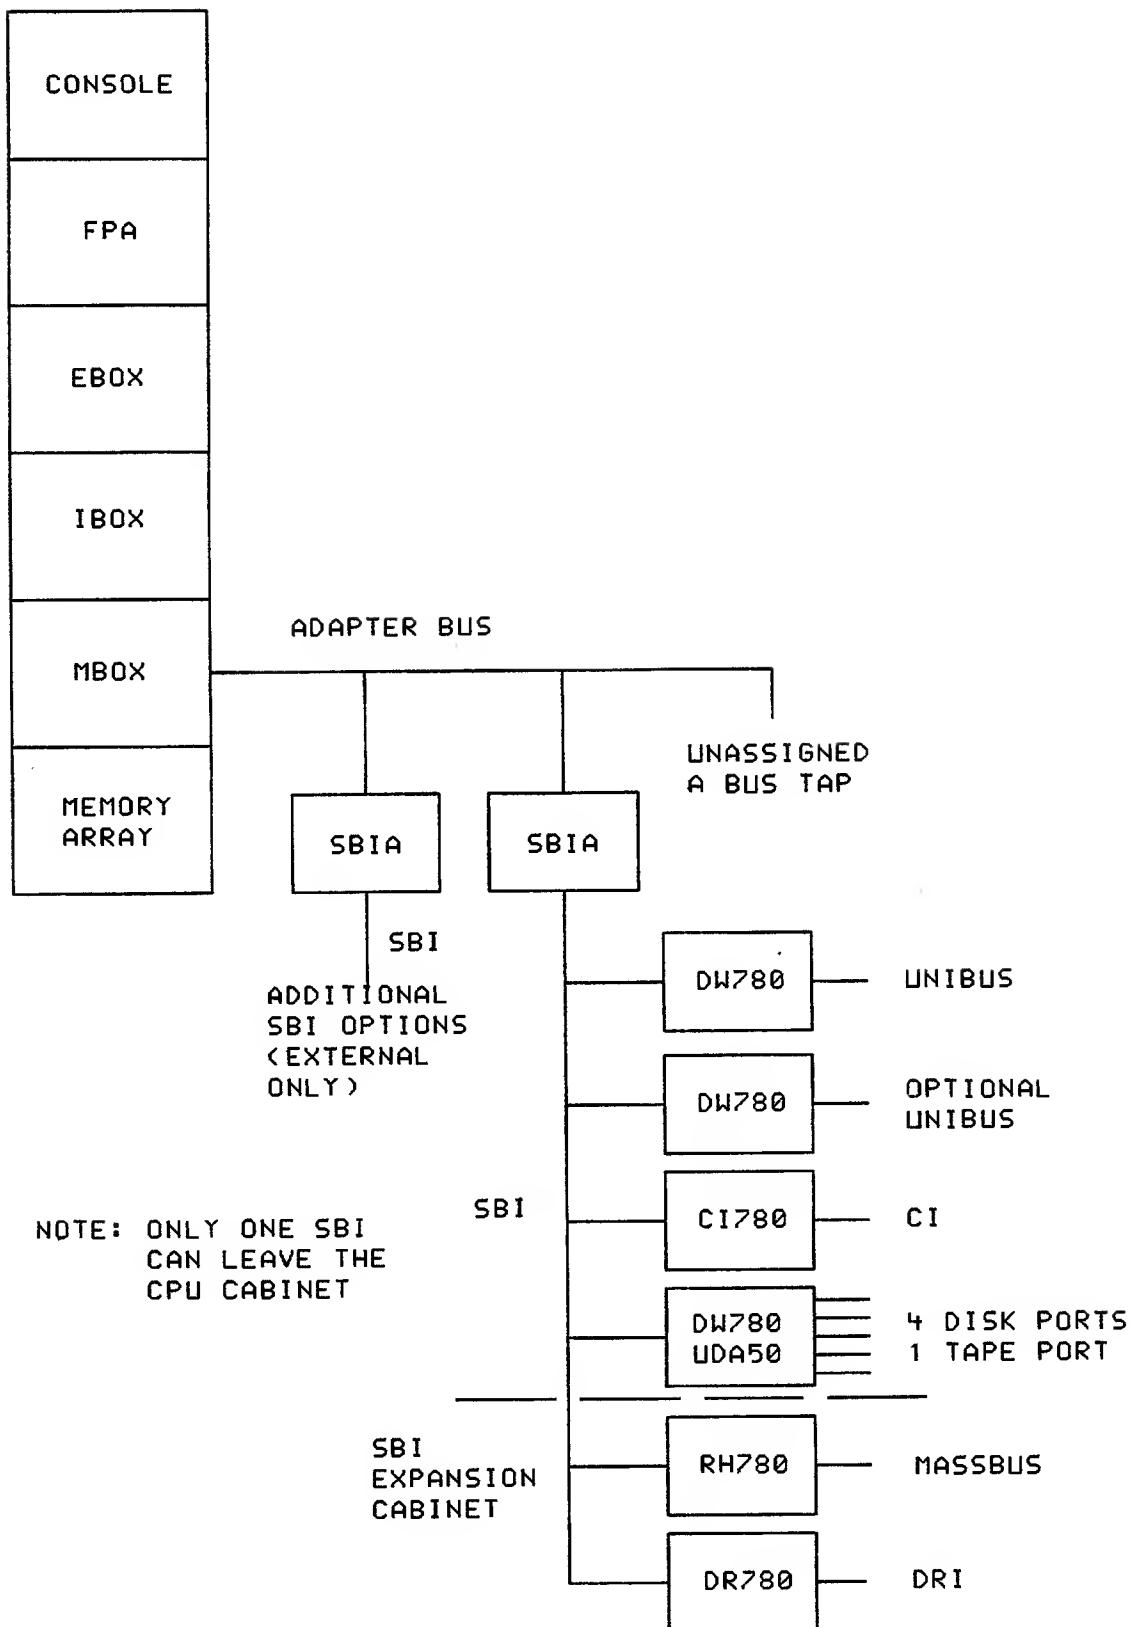


FIGURE 2.2 BASIC SYSTEM CONFIGURATION

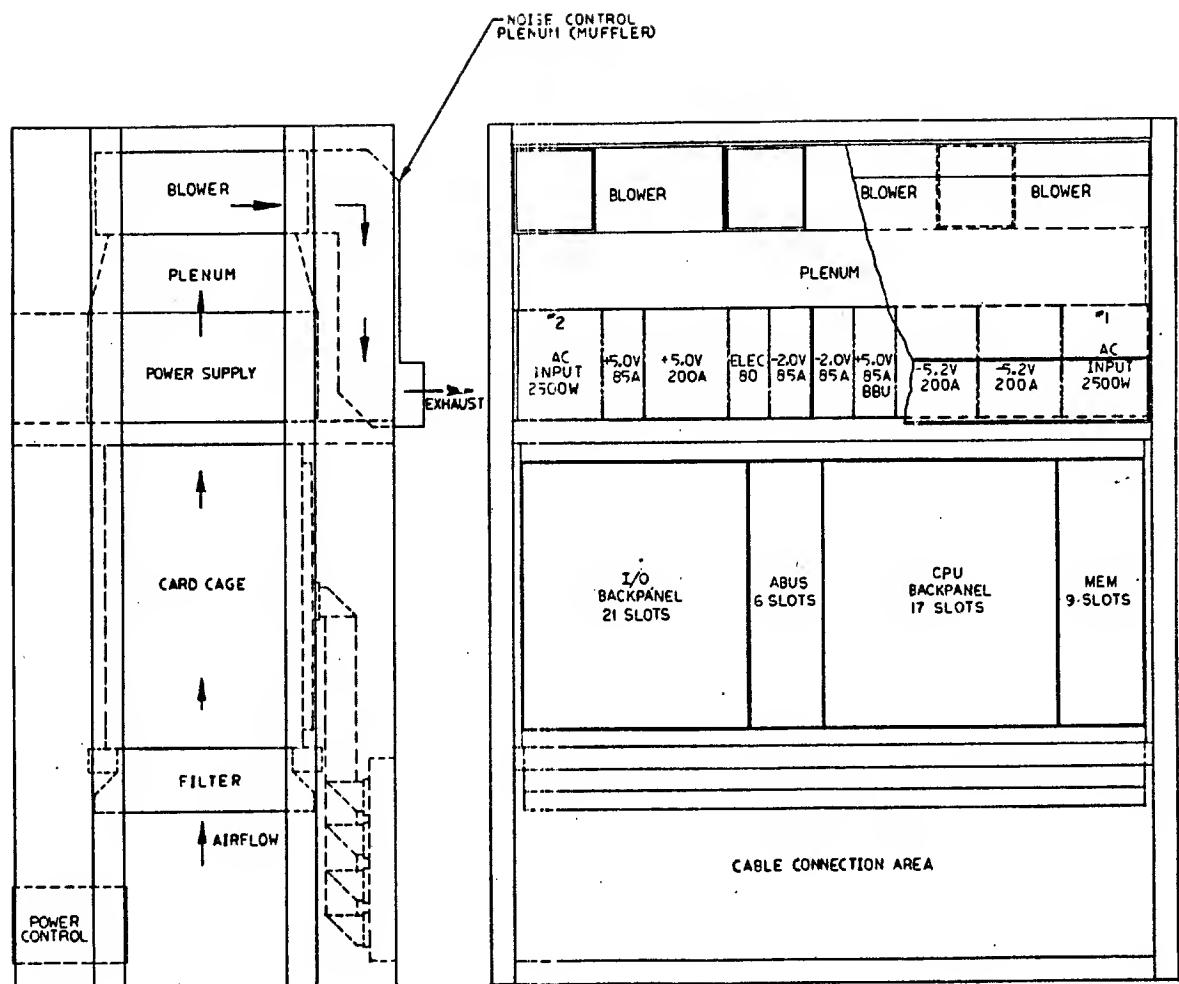
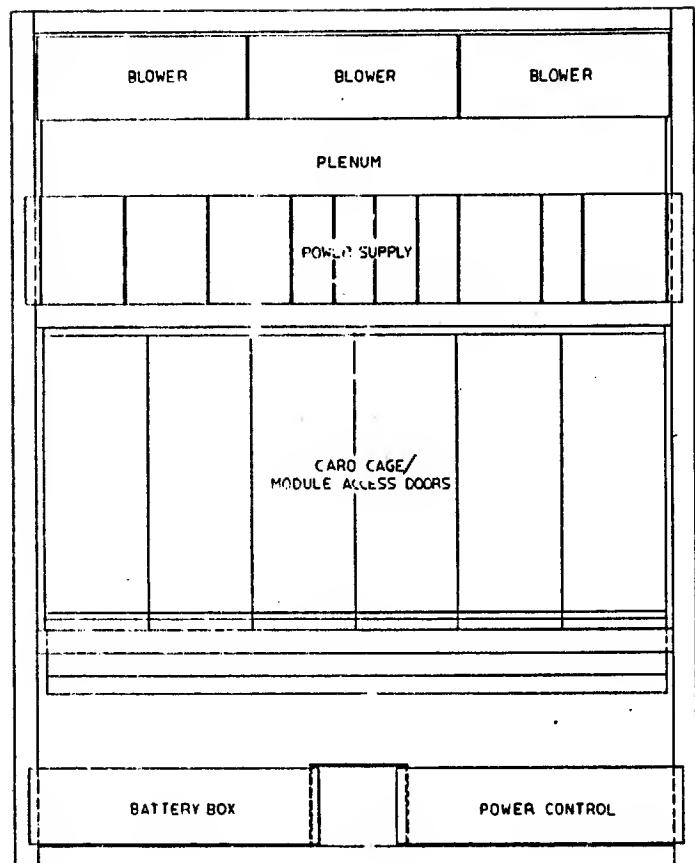


FIGURE 2.3 CPU CABINET CONFIGURATION

Modules with MCAs are one inch apart, those without use half-inch spacing.

The right half is also a single card cage containing the small A bus backplane and the larger IO backplane; this entire cage could be replaced for special applications or to make way for a mid-life kicker. The A bus backplane (on the left) has space for two SBI adapters of two L modules each (the first SBIA is plugged in at the farther end from the processor), and a spare slot for the third A bus tap. The IO backplane accommodates a mixture of hex, extended hex and L modules for the adapters on the internal SBI. This backplane contains one DW780, space for a second, optional DW780, and either a CI780 CI adapter or a DW780-UDA50 integrated disk-tape controller (although there is nothing to prevent a single system having both should that be desired). Any adapters beyond these must be mounted outside the CPU cabinet; such adapters can include the DR780 and RH780 (standard in Massbus systems) as well as those available in the CPU cabinet. In any system a single SBI can handle all adapters, or a second SBI can be added to handle the external ones (i.e. only one SBI can leave the CPU cabinet).

In systems built for shipment outside the United States and Canada, there will be a 10 kVA transformer to interface with the numerous variations in available utility voltages. This transformer is mounted at the bottom of the Unibus-console cabinet at the left of the CPU cabinet. Voltage conversion requirements for Unibus expansion cabinets, SBI expansion cabinets and the like will be handled in the fashion determined by the groups responsible for their design.

Cabinet Arrangements

Every system that has no more than 8 MB of memory has a Unibus-console cabinet attached to the left side of the CPU cabinet viewed from the front. Besides typical Unibus communication and unit record controllers, this cabinet (Figure 2.4) contains the console load device and the stepdown transformer wherever that is necessary. If additional Unibus cabinets are required, they are bolted on at the right. If external adapters are needed, they are installed in one or two SBI expansion cabinets (four per cabinet) attached at the right of the CPU cabinet, between it and any additional Unibus cabinets. Cabinets for Unibus disks and tapes are attached at the far right. CI and Massbus devices can be attached at the right or placed separately elsewhere in the computer room.

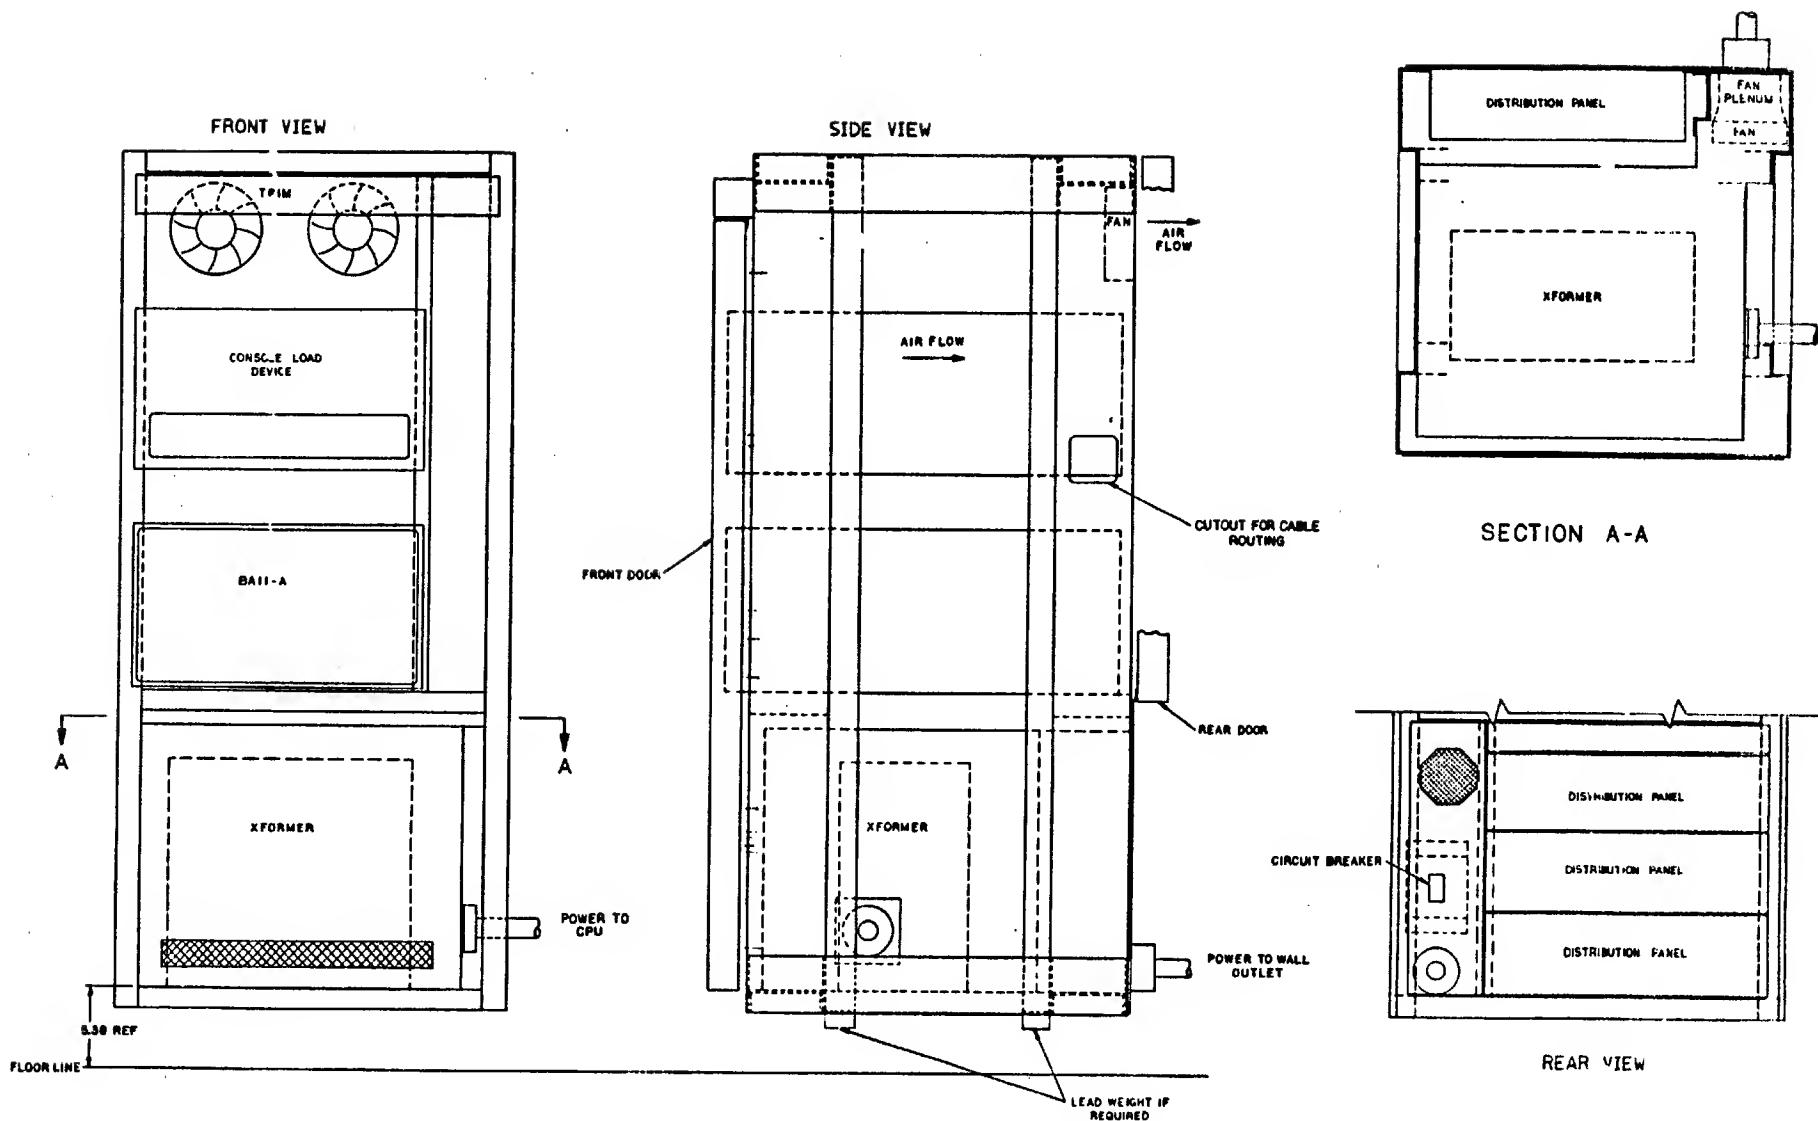


FIGURE 2.4 UNIBUS-CONSOLE CABINET CONFIGURATION

Should memory expansion beyond 8 MB be desired, a memory expansion cabinet (Figure 2.5) must be installed at the left between the CPU cabinet and the Unibus-console cabinet. This can accommodate an additional 24 memory boards bringing total memory capacity to 32 MB. It is expected that within the next few years, memory density will quadruple with the introduction of 256K MOS RAMs. At that time Venus will convert to 4 MB memory array boards, so the CPU cabinet can accommodate 32 MB of memory, and expansion to 128 MB will then be possible. Where required, a 5 kVa stepdown transformer for the additional memory is installed at the bottom of the cabinet.

2.3 PERIPHERAL EQUIPMENT

The peripheral equipment includes both new and current products in all categories: interconnects, adapters, controllers, and individual IO devices.

Interconnects

Venus configurations are based on system interconnects reflecting Corporate strategies. The characteristics and status of these interconnects are given below; all have high data integrity and provide appropriate electrical isolation.

- SBI Synchronous Backplane Interconnect ~ present VAX interconnect to Unibus, Massbus and other adapters; high bandwidth, medium cost.
- CI Computer Interconnect ~ joins loosely coupled computers, and mass storage, real time and communication subsystems; high bandwidth, 16 nodes; under development
- SI Storage Interconnect ~ joins disk and tape drives to controller; high bandwidth; being tested (the principal Venus version of the SI is the SDI ~ standard disk interface ~ and several disk projects are being based on it).
- NI Network Interconnect ~ joins computers, work stations, intelligent terminals, etc in local network; moderate bandwidth, moderate number of drops; currently being defined.

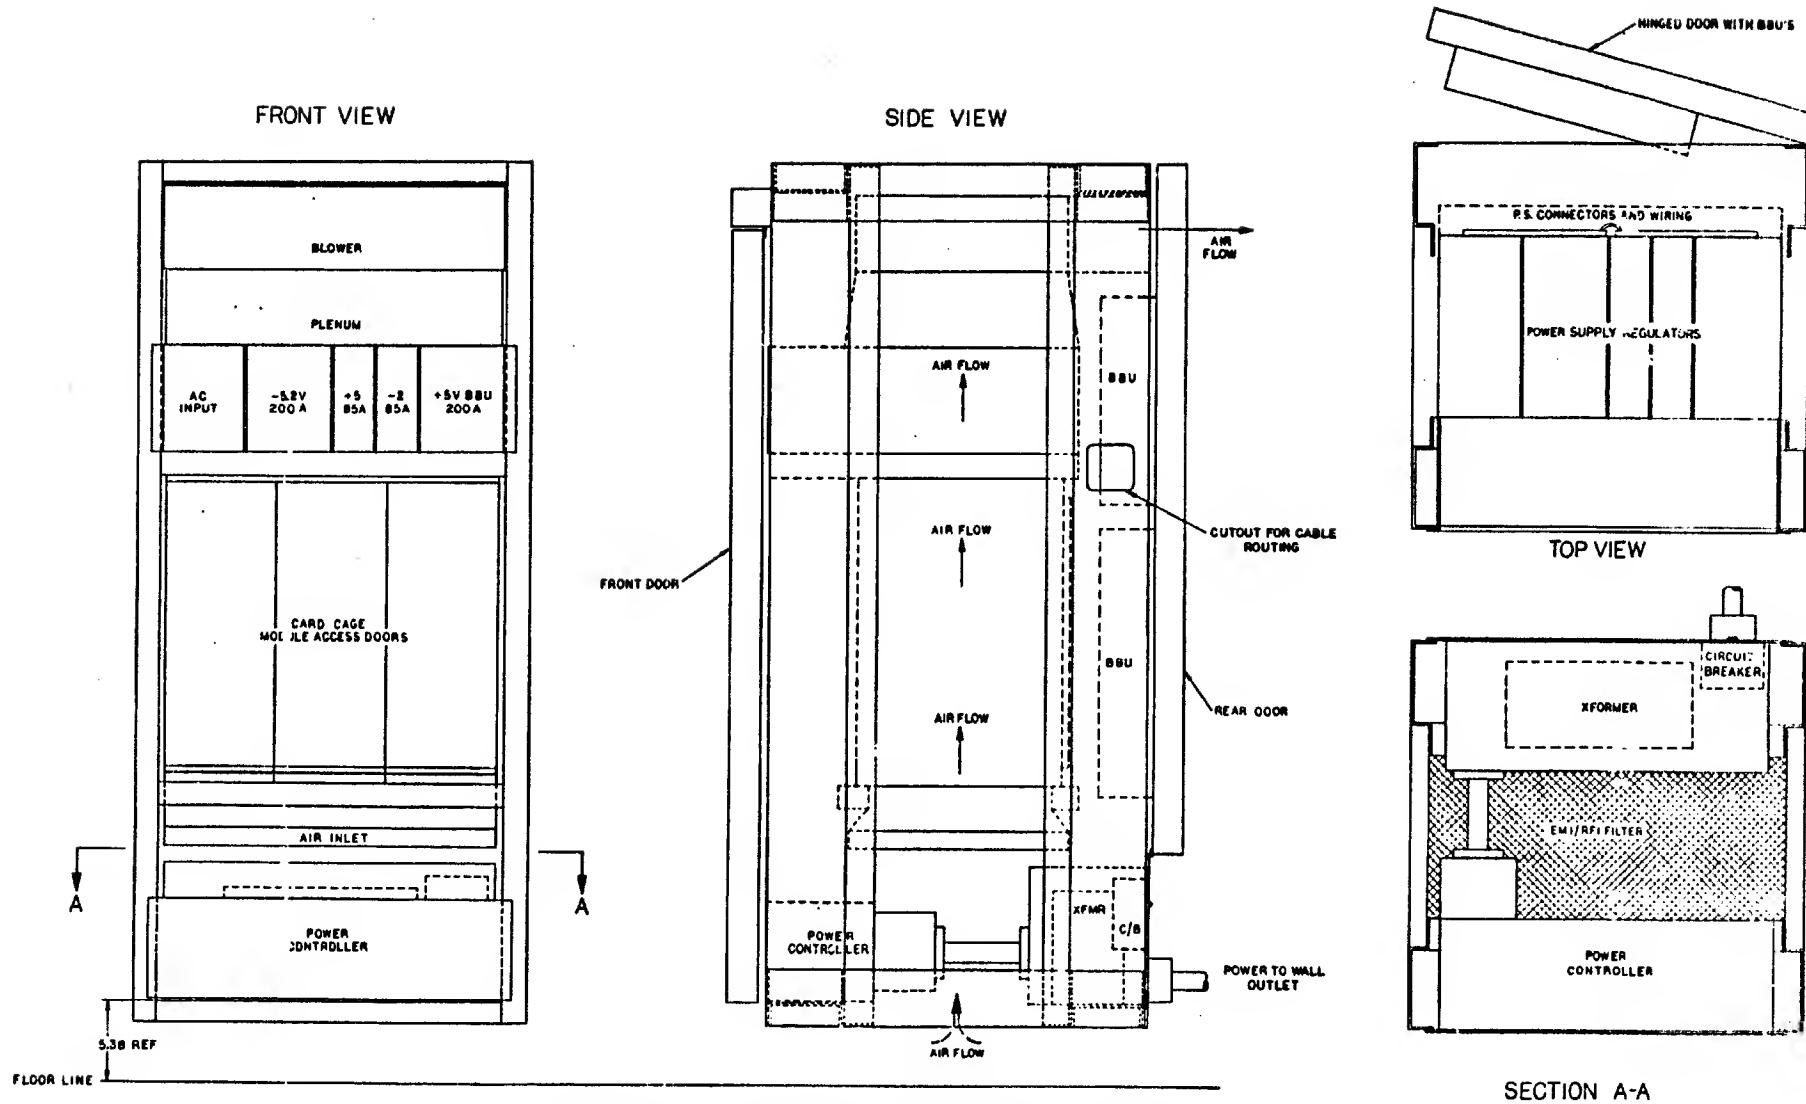


FIGURE 2.5 MEMORY EXPANSION CABINET CONFIGURATION

Adapters

These interface the A Bus to the SBI or the SBI to subsidiary interconnects such as the CI and Unibus.

SBIA SBI adapter on A bus, scheduled and funded.

DW780 Unibus Adapter - current product.

CI780 CI adapter, scheduled and funded.

DR780 High Speed Block Transfer Port for customer equipment or CPU-CPU communication; current product.

RH780 Massbus Adapter - current product.

Controllers

Two major controllers are expected to be utilized in the Venus system. Minor controllers, such as for unit record equipment, are included with the devices themselves.

HSC50 Hierarchical Storage Controller - intelligent mass storage controller on the CI; contains six subcontrollers for SDI disks and tapes; unit has many RAMP and high-performance features; project funded with FRS scheduled for Q4/FY83.

UDA50 Intelligent mass storage controller on the Unibus; contains four subcontrollers for SDI disks; some RAMP and high-performance features; project funded with FRS scheduled for Q2/FY82.

Recommended IO Devices

Following are those products we strongly feel should be developed for Venus. Current products we intend to support are listed at the end of Section 3.6.

Pinon removable disk on SDI, 180 MB capacity.

RA81 fixed disk on SDI, 400 MB capacity, 3 MB per second maximum transfer rate.

TA78 tape, 1600/6250 bpi, 125 ips, HSC50 controller.

LCGCR tape on Unibus, 1600/6250 bpi, 75 ips.

TU78 tape, 1600/6250 bpi, 125 ips, TM78 controller on Massbus.

2.4 SOFTWARE

As far as the software is concerned, the most important thing to keep in mind is that Venus is a VAX processor, that its operating system is VMS, and that there is only one version of VMS for all VAX processors. VMS development operates under two fundamental requirements: that no matter what processor the single version of VMS is booted on, it shall be capable of configuring the support for that processor automatically and transparently; and that a user program or higher level system program that runs on any VAX processor shall run on all. As shown by the examples in the chart below, the many elements that constitute the complete VAX software package are organized into three layers that are superimposed or built upon the VAX hardware.

Unbundled (Layered) Products	DBMS Fortran Coral-66	Datatrieve Basic Ada	Pearl Cobol APL	Bliss Pascal
Bundled Products	DECnet Linker Error Log Analysis	RMS-32	Batch Runtime Library	SYSGEN
Executive	Scheduler IO Adapter Support Processor Errors	Device Drivers Processor Initialization	System Services Memory Management	
			VAX Hardware	

In this structure, the bundled products are those programs that are intimately associated with the executive and that together with the executive constitute the basic software necessary for effective operation of the system. The higher level products in the unbundled layer (usually themselves referred to as the "layered" products) maintain a certain independence of the operating system, in that they are developed, released and sold separately. Such organization naturally requires that any changes in the lower layers be made in an upward compatible way. Unbundled products are not done directly by the VMS Group.

In terms of Venus software, it should be noted that processor dependencies are visible almost exclusively to the programs in the executive layer. Moreover any features or enhancements added to the layered products for Venus automatically become available to all VAX systems.

2.5 PERFORMANCE

The overall performance goal of the Venus Program is four times the performance of the 11/780.

Processor

As against the 11/780, pipelining in Venus halves the number of cycles required in most instructions, and the cycle time itself is reduced from 200 ns to 67 ns. To minimize processor idle time, the I box continuously prefetches the instruction stream, and the writeback cache eliminates 50% of the memory writes required. Taking all of these factors into account, the table on the next page compares our best estimate of instruction times on Venus against the times for identical instructions on 11/780. The figures listed are based on these assumptions.

11/780 main memory read access time, as seen from the E box, is 1400 ns. This is actually the best case, since it assumes the access is to the first longword, and the SBI and memory are idle.

Venus main memory read access time, as seen from the E box, is 533 ns. This does not account for writeback time or IO interference, which we believe to be negligible. Moreover it assumes the memory being used is in the CPU cabinet; read access to addon memory in an expansion cabinet takes 25% longer (overall system performance degradation with expanded memory is 5-10%).

11/780 write cycle time is 200 ns. This does not account for contention at the SBI write buffer or in the memory control.

Venus cache write cycle time is 67 ns. This does not account for writeback, and assumes aligned longword writes. Masked writes take two cycles.

There is no contention for general register access while waiting for updates. Such contention would occur primarily when a register destination of one instruction is used in the address calculation of the first operand of the next.

Memory references are aligned.

Operation	100% cache hit 11/780	Venus	Memory reads/ writes	Ratio at hit rate		
				100%	95%	90%
ADDL2 R,R	400	67	1/0	6.0	5.2	4.8
ADDL2 ^B(R),R	800	133	2/0	6.0	5.2	4.9
ADDL2 R,^B(R)	1200	267	2/1	4.5	3.9	3.6
ADDL2 ^B(R),^B(R)	1400	267	3/1	5.2	4.4	4.0
MCVL R,^B(R)	800	200	1/1	4.0	4.6	4.0
MOVL ^B(R),^B(R)	1000	200	2/1	5.0	4.1	3.7
INCL R	400	67	1/0	6.0	5.2	4.8
INCL ^B(R)	1000	200	2/1	5.0	4.1	3.7
ADDL3 R,R,R	600	133	1/0	4.5	4.3	4.2
ADDL3 ^B(R),^B(R),R	1200	133	4/0	9.0	5.1	4.7
ADDL3 ^B(R),^B(R),^B(R)	1600	267	5/1	6.0	4.7	4.3
BR successful	600	200	1/0	3.0	3.1	3.2
BR unsuccessful	400	133	1/0	3.0	3.1	3.2
MOVC3 per byte (1)	300	33		9.0		6.0
MOVC5 clear per byte (2)	300	16		18		5.5
ADDF2 R,R	800	133	1/0	6.0	5.5	5.3
Small difference	800	200	1/0	4.0	3.9	3.9
ADDF2 ^B(R),R	1400	133	2/0	10.5	8.2	7.1
Small difference	1400	200	2/0	7.0	6.2	5.7
ADDF2 ^B(R),^B(R)	1800	267	3/1	6.7	5.4	4.8
Small difference	1800	333	3/1	5.4	4.6	4.2
ADDD2 R,R	1400	333	1/0	4.2	4.1	4.1
Small difference	1400	400	1/0	3.5	3.5	3.5
ADDD2 ^B(R),R	2200	333	3/0	6.6	5.9	5.5
Small difference	2200	400	3/0	5.5	5.1	4.9
ADDD2 ^B(R),^B(R)	3600	533	5/2	6.7	5.5	4.9
Small difference	3600	600	5/2	6.0	5.1	4.5
MULF2 R,R	1200	200	1/0	6.0	5.7	5.4
MULF2 ^B(R),R	1800	200	2/0	9.0	7.3	6.9
MULF2 ^B(R),^B(R)	2200	333	3/1	6.6	5.5	4.9
MULD2 R,R	3200	733	1/0	4.4	4.3	4.3
MULD2 ^B(R),R	4600	733	3/0	6.3	6.0	5.7
MULD2 ^B(R),^B(R)	6200	867	5/2	7.2	6.2	5.6

Notes:

- (1) MOVC3 times in the 90% column assume steady-state cache performance for long strings - the 11/780 waits for memory read every eighth byte, memory write every fourth; Venus waits for read and write every sixteenth byte.
- (2) MOVC5 clear times in the 90% column assume steady-state performance for long strings - the 11/780 waits for full memory modify cycle on each longword and Venus waits for writeback every sixteen bytes.

Venus cache miss ratio is assumed to be half that of 11/780, because the Venus cache is twice as large. Thus the data shown as 95% hit rate compares Venus at 95% with 11/780 at 90%, and the 90% column compares Venus at 90% with 11/780 at 80%.

100% cache hit times were calculated by counting the maximum number of cycles required by the given instruction from each box, M, I and E, and multiplying by 67 ns. 95% and 90% times were then obtained by adding an "average miss penalty" for each memory access performed by the instruction. For 11/780, this penalty per location read is 120 ns for the 95% hit rate, 240 ns for 90%; for Venus it is estimated at 33 ns for 95% and 65 ns for 90% per location read or written.

In Venus, 3-operand instructions with register destination take one cycle longer than the 2-operand equivalents with register destination. With memory destination, 3-operand instructions take the same time as the 2-operand forms, unless the destination specifier requires indexing or indirection.

Floating Point Performance

The lower half of the table shows the most recent attempt to quantify the performance of the Venus floating point accelerator as against that of the 11/780. There are still many uncertainties in these estimates:

The single precision add path is extremely tight. The large effect of going from two to three cycles is a high incentive for making it work, but this is still a major risk.

The microsequencer and control store in the FPA have not received significant attention. Neither is expected to present a major problem, but we need to be sure.

It may be feasible to gain significant improvement in the multiply times by enlarging the multiplier, but this is a complex tradeoff involving the size of the control store, the clocking scheme, and the pinouts available on the adder module. We have therefore set expectations on the basis of a 16 x 16 network.

Note that two sets of statistics are given for all additive operations, the second being for an ADD or SUB involving a small difference in the operands,

i.e. an operation in which one fraction is subtracted (add where the signs differ, subtract where the signs are equal) and the exponents differ by 0 or 1. This is the case where normalization may require a long shift or negation of the fraction, which in the Venus FPA takes an extra cycle.

The comparisons above are of course with FPAs installed in both 11/780 and Venus. For machines without FPAs the improvement is not as good, but it is assumed that any customer concerned about floating point performance will buy an FPA; it is therefore not worth the cost to make any special effort to improve non-FPA performance.

Memory

Not only has the memory cycle time been reduced considerably, but once a write cycle has been started in one array board, the M Box can initiate a second write in another. Timing is such that different array boards can absorb writes as fast as the M Box can request them. Memory access times in nanoseconds at the M Box are as follows (a block is four longwords, sixteen bytes).

	In memory	From CPU
Read access, first word		
In CPU cabinet	400	533
In expansion cabinet	533	667
Block read access		
In CPU cabinet	600	733
In expansion cabinet	733	877
Block write access	267	
Block write cycle	467	
Overlapped operation		
Block write access	267	
Block write cycle	333	

Figures from CPU include time for request from CPU to memory and time to get data from memory back to CPU. No write times are given for the CPU, as writing in main memory is only for writeback; if this delays the CPU, the wait depends on a variety of circumstances. The write cycle time in overlapped operation is the time between write requests to different array boards.

Peripheral Subsystem

The overall performance of the IO subsystem depends both on the performance of its individual parts, and also on the characteristics of processor and memory, as the peripheral adapters and controllers interact so

much with them. For a typical mix of traffic, the memory burst bandwidth on the A bus is 13.3 MB per second with one SBIA, 17.1 with two (for details refer to Appendix B in the IO Plan). The frequency of interrupt checking will be greater than in the 11/780, reducing both the latency and the worst-case wait. The following table gives the maximum throughput or data transfer rates in megabytes per second for the various units in the IO subsystem.

SBI	13.3
Unibus	1.5
CI	5.25
DR780	6
Massbus	2.2
SI	3.125
HSC50	6.25
Disk	3.125
Tape	2
UDA50	1.5

Maximum A bus bandwidth with two or more ultra high speed adapters using octaword transfers would be 34.3 MB per second.

Software

Software performance depends not only on the activities of those who create the software, but also on the performance of the hardware on which it runs. Hence the most direct improvement of VMS as an operating system for Venus will be the improvements in processor speed, memory size and speed, and IO capabilities of Venus as against the 11/780. On the other hand, each release of VMS has goals for quality, functionality and performance, among others, to varying degrees. Venus will usually benefit from performance enhancements engineered for the whole VAX family.

2.6 RAMP AND DATA INTEGRITY FEATURES

Features to guarantee the integrity of the data in the system and to promote its reliability, availability and maintainability are built into Venus at every level: they range from minor characteristics of individual circuits to major provisions embracing the entire system. Some of the more significant features are these.

Inherent reliability achieved through low component count and worst-case logic design.

Dynamic monitor error reporting, by means of an error logger, to aid in identifying the source of an intermittent failure. This logger will be used for both hardware and software malfunctions. The log is kept in a disk file.

Instruction retry whenever it is appropriate to the error type. For example five copies are kept of the general purpose registers. Hence on a GPR parity error, the instruction can be repeated using a different copy.

Additional software features in this area include automated patching and updating procedures, powerfail-restart support, user mode diagnostics, extensive protection facilities, and dynamic memory configuration to exclude bad pages.

Parity checking at all RAMs and buses, and parity continuity carried through all major data paths. Parity is kept not only for data, but also for physical addresses and the microcode.

Separate selects to each memory array board, so the control logic for storage selection is all in one place, and faults can be isolated to an individual board.

Single bit error correction and double bit error detection for the cache and memory array, with automatic rewriting of the corrected word.

Memory battery backup for 10 minutes. Backup can be set shorter to save on battery recharge time, thus allowing user to choose riding out several short power failures at the cost of going down during a long one.

The ability to reconfigure the system without the FPA when troubleshooting floating point failures.

Optional redundant IO transfer paths. For example, interconnected disk systems, so that should one controller fail, an individual disk can maintain a transfer path through another (the switchover must be handled by the operator however).

Fast, accurate diagnostics, with first-failure fault isolation to field replaceable unit in CPU cluster and to module in peripherals.

Error logic for monitoring all backplane signals from the console via the diagnostic bus.

A "keep alive" count kept by the console to determine if the system is hung. Should the hung condition be detected, the console saves the state of the machine.

Other console diagnostic features, including remote capability, flexible clock control, and some visual indicators should the console be unable to report its own failures.

For high availability, the Hydra configuration, where one processor takes over all activity should the other fail. The three criteria for Venus to qualify as a Hydra node are: it runs on VMS, it has a CI, and the console uses the standard midrange console language.

Extensive console monitoring of the environment and the power system.

3 DEVELOPING THE SYSTEM

The Venus Program is organized around a strategy for developing a complete computer system. This strategy has a number of stages, and it is accompanied by phase reviews and thorough program tracking, with major milestones pinpointed. The process culminates in the marketing of two classes of systems at specific times. There are of course risks attendant on the development process, and these are duly noted.

Almost all of the staff and all key personnel are already in place. Figure 3.1 illustrates the Program organization.

3.1 DEVELOPMENT STRATEGY

The overall strategy is to develop a system whose IO subsystem is based on the SBI, using the CI for mass storage and the Unibus for communications and unit record. However Massbus systems will be available for backup, and employment of a Unibus-UDA50 integrated disk-tape controller in place of the CI will allow systems at lower cost with moderate IO performance. The reasons for taking this approach are several:

SBI and the Unibus form a known, well-defined bus structure with mature software, and the CI with its mass storage peripherals is well enough along so we can depend on it with confidence.

Improved cost/performance ratio of CI-HSC over Massbus in medium-to-large system configurations.

Better chance of meeting FCC and acoustic goals with CI peripherals.

Satisfy current 11/780 customers via upgrading.

Offer Unibus and Massbus upgrades.

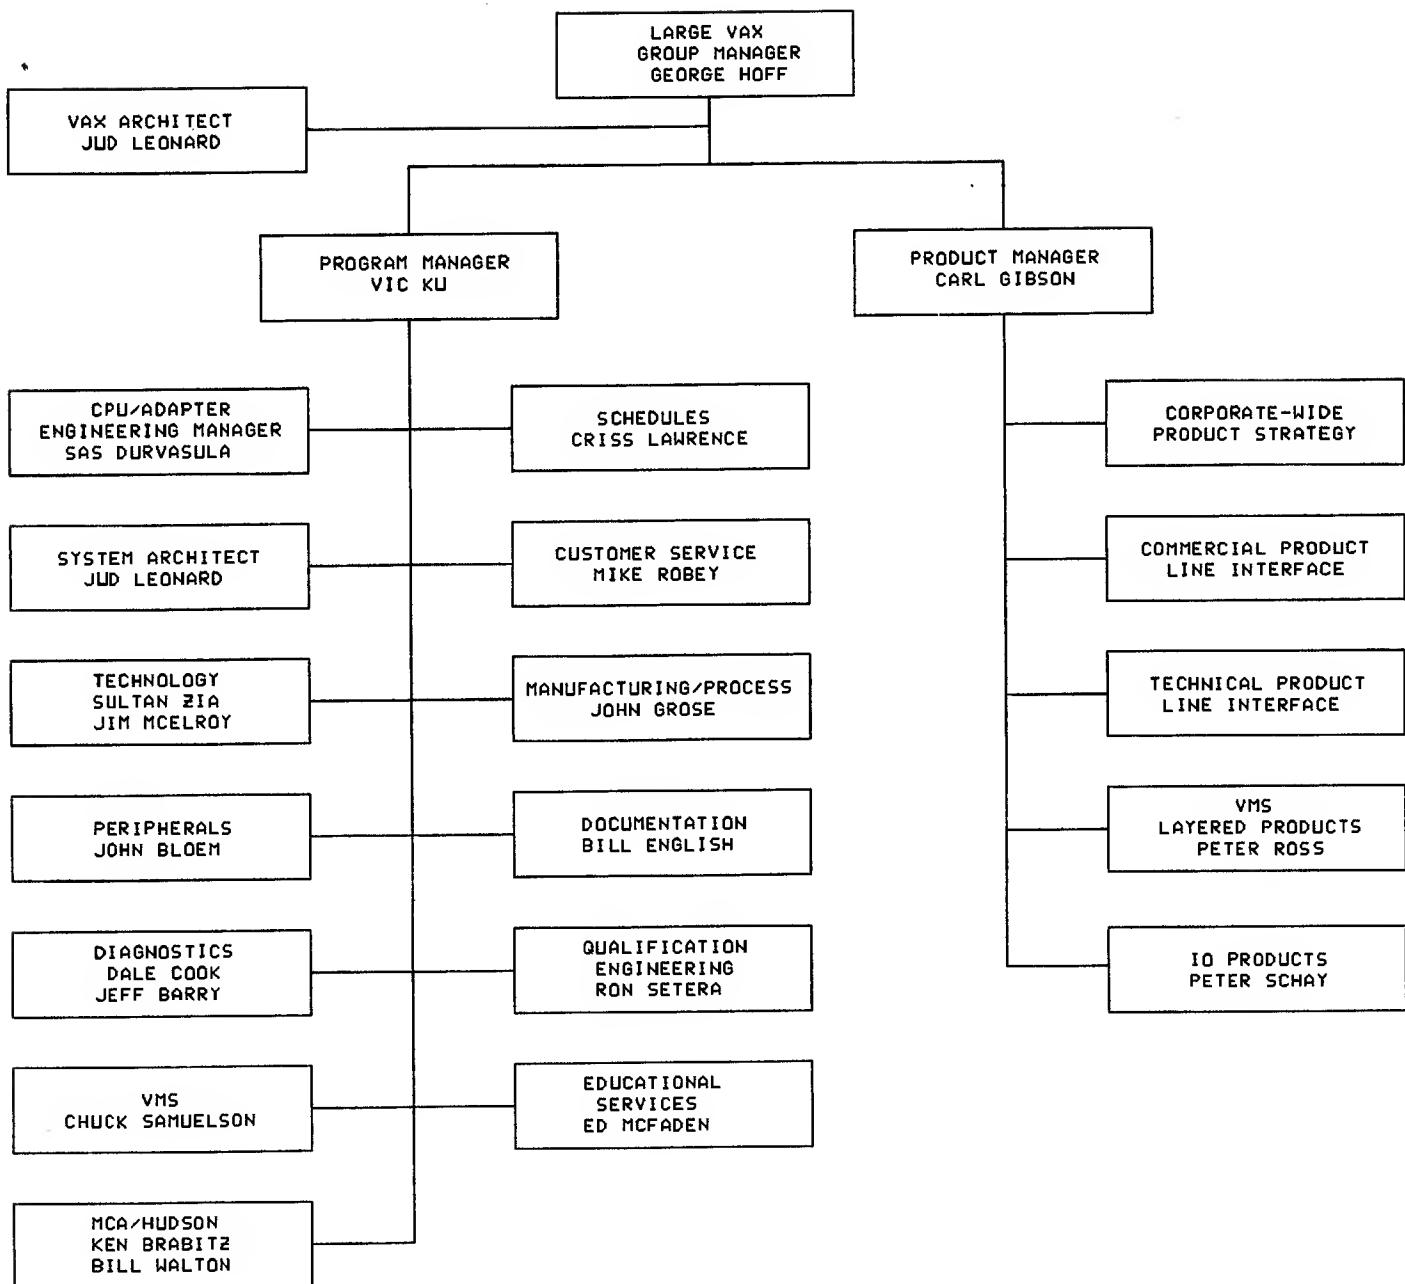


FIGURE 3.1 VENUS PROGRAM ORGANIZATION

The Venus development strategy is organized in the following seven stages.

1. System Performance Analysis

The Venus Program is funding two people in the SPA Group in Maynard to do system and IO performance analysis and system simulation to determine whether the Venus structure is properly designed conceptually. They will use benchmarks and workloads to determine whether functionalities of the different parts of the system are consistent with one another, e.g. are there any bottlenecks? is there enough IO bandwidth?

2. SAGE2 Simulation

Throughout the Program, before any element is actually built it undergoes complete SAGE2 simulation. This occurs at both the chip level and the box level (e.g. I and E boxes); simulation and delay analysis at the complete system level will be started well ahead of breadboard power on.

3. Technology Evaluation

Concurrent with the above two stages, the Technology Group has been evaluating the various technological innovations that are under consideration for use in Venus. This involves generating specifications, working with vendors to meet those specifications, and collecting the information on which the Program can base its decisions on which technologies to use and what tradeoffs to make. The most extensive evaluation procedure is that for the MCAs, discussed below.

4. Breadboard Stage

Engineering will build two breadboards that will run at reduced speed but will allow verification of system functionality. It is expected that the functionality of the breadboard will be equivalent to that of the final machine, but revisions will be made as necessary. This stage in the development will enable Diagnostics and Software to begin debug. It is planned that people from Manufacturing, Customer Service, and System Qualification will participate; they will begin learning the system in terms relevant to their own tasks, and will be actively involved in identifying problem areas where Engineering can help them. Moreover since a simulator is not a real machine and also runs perhaps a million times slower than a real machine, the breadboards will

allow real-world functional verification of the simulated hardware as well as procedures that could not be simulated at all.

The breadboards will use backplanes that are mostly wirewrapped and printed circuit modules. We are pursuing an effective emulator strategy in case we require individual hardware simulators for any MCAs that are not functional. It is expected that at least 90% of first-pass MCAs will be available and 30% of them will be operational at this stage. Emulators will automatically be built for the thirteen most complex MCAs.

5. Post Breadboard and Prototype Stage

Engineering will build and drive five machines that implement a transition from advanced breadboards to prototypes. These will run at full speed, will use all etched pc boards and backplanes, and will make use of the second-pass to third-pass MCA chips. The use of these five machines is: one for preliminary 102 testing, one for preliminary DMT, one for software development and VMS verification, and two for engineering and diagnostic operations, system quality verification, and the like. People from Manufacturing, Customer Service, and System Qualification are expected to be deeply involved in these activities.

6. Manufacturing Prototype Stage

Manufacturing will build twenty-five prototypes and will then update them to reflect the documentation when it is released. The goal is for these to be functionally identical to the Engineering prototypes. The Manufacturing prototypes will be built with Engineering support and maintained by Customer Service. These machines will utilize the revised versions of the pc boards and backplanes, and they will have the second- and then third-pass MCA chips. Checkout will begin when the Engineering prototype runs AXE and the diagnostics. Disposition of these machines will be:

One to the Memory Group to test array modules,
One to Production for FA&T,
Four for final DMT verification,
Two for Field Service training,
Six to Production for quick verification,
One to Spitt Brook for the software team,
One for System Evaluation,
One for Mass Storage Production, and
Eight will be assigned by the Product Manager.

At this stage, System Engineering will use the prototypes to verify a limited number of system configurations. Since the initial prototype systems will be built and verified without released documents, we will set up an account for purchasing components and mechanical parts and to track actual costs instead of standardized costs.

7. Advanced Test Stage

This final stage is for testing and verification (qualification) of many more system configurations utilizing CI and Unibus devices.

MCA Engineering and Evaluation

The engineering and evaluation effort in MCA technology lies in three principal areas: development of software design tools, verification of the hardware design, and characterization and qualification of the finished parts.

The various steps in the creation of an MCA use many of the design, layout and testing tools discussed in the next section. The most important new tool in layout is MCACUT, the MCA version of the Merlin placement optimization system. This uses a cutline approach that minimizes the number of nets crossing an imposed boundary by swapping equivalent entities; these entities may be complete cells, equivalent functions within cells, equivalent gates within functions, and equivalent pins within gates. Other parts of the project include enhancements to the IDEA circuit routing system, programs for creating and checking the artwork data base that serves as the input for making the IC mask, and programs for manipulating and analyzing test patterns.

The hardware part of the project includes determining those circuit parameters that predict operational performance of the chips, to verify that the various circuit elements implemented on a chip function correctly - both in theory through circuit simulation and in practice through evaluation and test verification - and to determine worst-case conditions that may be applied to automatic test hardware to support verification. The hardware development group also provides the necessary application support to the Venus design team and serves as the technical interface to Motorola.

Completing the development strategy is a complete characterization of the test chips, followed by a qualification procedure that includes correlation among test systems, compliance with specifications, verification of the burn-in procedure, verification of performance of parts from several wafer lots, and thermal, mechanical, electrical and reliability testing.

Related functions are installation and organization of MCA manufacturing and test procedures in the Hudson facility and acquisition of parts.

Software Strategy

To ensure an operating system for Venus in the hardware development time frame, the development strategy of the VMS Group in Spit Brook is to be heavily involved as an integral part of the Venus team throughout the design of the system. The following strategy is based on the belief that there are no areas where the VMS executive lacks the mechanisms and structure to support the changes required by Venus; it is expected that enhancements to the current system will be sufficient for Venus support.

1. Place product quality as the primary and overriding goal above addition of new features.
2. Have a VMS release with appropriately tested Venus support in the SDC when Venus ships, with a fallback strategy of releasing VMS with partially tested support and providing binary update patches when final testing is accomplished.
3. Assume that VMS does nothing to interfere with or bottleneck the system performance for a Venus class processor. This assumption will be tested empirically by the System Performance Analysis Group. If the assumption proves false, the scope of this project will be modified; the answer to this question must be determined before the Phase 1 review for the VMS release that supports Venus.
4. Schedule and implement VMS support for new-interconnect adapters and peripherals as they become available. The general rule of thumb is that to support a given peripheral, firm specifications must be available one year before support is to be provided, and hardware must be available to VMS Development six months before submission of the VMS release to the SDC. However exceptional devices such as HSC50 will require longer leadtime.

5. The present Venus-VMS Project is restricted to activities in the executive layer of the software in three categories: booting and system initialization, system error handling, and IO. Venus requirements that imply software development in the bundled products will be scheduled and implemented by the VMS Group in the normal phase review process for the appropriate VMS release. Even though the current project will not directly result in the implementation of bundled products, a major role of this project is to provide a communication and consulting path between Venus and VMS. As Venus product requirements are identified, this project will provide a focus to help ensure schedule and technical needs are made known and satisfied within the constraints of the phase review process.
6. Venus requires the development of layered products. Layered product requirements are defined by the Venus Program, in which this project is a participant. The products themselves are defined and implemented using the phase review process, managed by the responsible development group. Conformance to the appropriate VMS standards and testing and integrating them into VMS is monitored by VMS program management. Refer to Appendix C for the current status of the layered products.

3.2 DEVELOPMENT TOOLS

Being developed internally are a number of programs, many very large and complicated, to aid in hardware design. Although many of these tools are available throughout the Corporation, most are being developed or enhanced specifically for the Venus Program by the LSG CAD Group. These programs are generally referred to as CAD tools, for "computer aided design".

Basic Circuit Design

The fundamental CAD tool is SUDS, the Stanford University Design System. This is based on a sophisticated graphics editor that aids in the design and checking of logic circuits and drawing of circuit schematics. It also contains programs to create wirelist and plot files. The outputs of SUDS provide the inputs to CALDEC, IDEA, and most of the other software discussed below. The program also has facilities for interacting with the various intermediate products of the board and MCA layout procedures.

SAGE2 Simulator

From information provided by the SUDS wirelist file, SAGE2 simulates the hardware of individual MCAs and Venus subsystems with the ability to inspect the interaction between individual gates in real time (although many times slower than actual gate speeds), to determine whether the logic actually does what it was designed to do. The whole system can also be simulated, with inspection at levels higher than individual gates.

VOTE Simulator

This program determines the effectiveness of the test patterns, test programs and test microcode generated by Diagnostics and Manufacturing. VOTE effectively runs the tests on a simulation of the hardware with the equivalent of physical fault insertion done entirely in software, and from this determines which of the faults the test was unable to detect.

IDEA

This program is the successor to CALDEC. Like the earlier procedure, it uses SUDS outputs to lay out circuits, but it is more advanced and handles many more layers. Only IDEA has the capability needed for laying out MCA chips.

Delay Calculation

This software package allows the designer to determine the physical delays between individual signal points in a circuit design, either a single board or a set of boards in a backplane. From the SUDS wirelist file and the CALDEC output (eventually the IDEA output), DLY creates a database that represents the physical hardware as it would be built, and from that CAL calculates all signal propagation times taking into account gate delays, wire links, and even stubs. Then with DLYED, the user can determine the delay structure of his design by inspection of propagation times across individual elements in each signal path. For MCA inspection, a file equivalent to the CAL output can be generated directly by SUDS.

Merlin Placement Optimization System

These programs help the designer determine the optimal position of circuit elements from critical parameters supplied by the designer and known characteristics of the materials, including even the capacitance of metal runs. The original program (MINCUT) is being enhanced in two stages to handle first MCAs and then pc and

multiwire boards. With the information provided by this software, the designer can go back via SUDS to SAGE2 to get real delays.

MCA Verification

From the IDEA database, the TENART software checks design rules, verifies interconnections, and ultimately generates the CALMA database, which is the representation of the MCA design used by Motorola to create the chip.

Wirewrap

From wiring rules and from information about the special character of runs, their type and termination supplied by the user, the wirewrap package generates the pattern for wirewrapping a circuit board or backplane, including assigning twisted pair grounds and generating an NC tape.

Test Pattern Generation

The Digitest D-LASAR program implements an algorithm that generates input and output test patterns for both simulators and hardware testers. This is used principally for MCA designs, which are generally sent to Digitest, but arrangements have been made to use the program in-house.

3.3 PHASE REVIEW

The Venus Program will follow the "Product Development Process" used by the Large System Group in Marlboro. This document details the entry and exit criteria for six phases, listed here with their completion dates. (Copies are available from Terry Mahoney, MRL-2/E78, DTN 231-6270.)

Phase 0	Product Strategy and Requirements	Completed Q2/FY80
Phase 1	Product Definition and Planning	Completed Q2/FY81
Phase 2	Product Implementation	Q3/FY83
Phase 3	Product Qualification, Product Release, and Pilot Production	Q4/FY83
Phase 4	Product Continuation	
Phase 5	Product Retirement	

During Phase 1 the many detailed plans and specifications were created, the design reviews held, the necessary contracts signed, and the manufacturing plant selected. The culmination of this phase was the completion of the Product Business Plan and System Implementation Plan, and the signing of the Product Contract. These last three items, which constituted the subject of the Phase 1 Review, detail the course of the Program in Phase 2. During Phase 2 the Program will accomplish the following major tasks.

Renew the commitment of Engineering, Manufacturing, Marketing and Customer Service to the revised Product Business Plan.

Hold technical reviews of the engineering design, manufacturing process, and service process.

Complete the detailed design, hardware breadboard and prototype test, and software internal tests; at the end of Phase 2, DMT and field testing will be ready to begin.

Run performance benchmarks and publish a performance report.

Make sure announcement criteria can be met. The principal criteria are that DMT be one-third done and there be a reliable second source for MCAs.

3.4 PROGRAM TRACKING AND MAJOR MILESTONES

The phase review process discussed above defines a number of major stages covering the entire history of a product. But to be able to determine the exact status of the Program at any time, and thus to know whether it is on target and to identify problem areas for taking corrective action, the entire development Program - at several levels - is continually undergoing an exhaustive and exacting tracking procedure employing several mechanisms.

The main tracking procedure is carried out at three levels, employing the PANTT Program Management System, which generates pert and GANTT charts. At the bottom level the individual project engineer uses a pert for keeping track of the detailed day-to-day activities of his part of the Program (such as mechanical packaging, MCAs, mass storage, etc). Each project engineer also works with a dedicated Program Scheduler at the intermediate level to prepare and maintain an overview schedule of the individual area and to determine major milestones and dependencies.

This level also makes use of waterfall charts for keeping track of the actual times that milestones occur as against their projected times. At the top level is a complete system overview that is maintained by the Scheduler from the information supplied by the individual area overviews. The Scheduler keeps a detailed top level pert and uses PANTT for a continuous analysis of the scheduling interdependencies among the various areas. This provides a focused look at each area from above, and allows the Program Manager to know on a day-to-day basis how a delay in any area will affect the others. The advantage of using PANTT is that whenever any change is entered at any point, its effects ripple throughout the entire structure so that one can tell immediately what those effects are on every other part of the Program. To do this manually would be impossible.

Since by far the largest category of operations in the entire Program involves layout procedures - for MCAs, pc boards, backpanels - another Scheduler uses PAC II to track just this single category. Overall the category involves three major stages: the creation of the MCAs, which are then combined on pc boards, which are in turn brought together at the backpanel. Moreover each of the stages is itself a multistage process. As an example the creation of an MCA requires these five steps:

1. Engineering and simulation
2. Layout
3. Verification
4. Fabrication
5. Test

In each step of the process a coordinator has responsibility for internal operational flow, and these coordinators work with the Scheduler to establish the entry/exit criteria for the various steps, i.e. the conditions which determine when an MCA design is ready to move from one step on to the next.

Within the Program are other tracking mechanisms, in particular a regular schedule review and an actively kept loose-end list. External to the Program per se are mechanisms effected by groups that operate in parallel with and provide services for the Program. Such tracking occurs mostly through site managers, such as those for Educational Services (particularly documentation) and in the service/process area (service includes diagnostics and drafting).

Major Milestones

Throughout the course of the Program, there will be a number of particular events or milestones that must occur at specific times if the overall goals of the Program are to be met. Such milestones are here given for five areas: the overall program, technology, the processor, peripherals and software. Keeping track of the actual dates of these events as against the target dates given here will provide a very good indication of whether the Program is on schedule and where any trouble may lie.

Program

End of Phase 1	12/80
Engineering breadboard power on	7/81
Engineering breadboard runs diagnostics	9/81
Engineering breadboard runs VMS	1/82
Engineering post breadboard power on	3/82
Engineering post breadboard runs VMS	5/82
Engineering prototype power on (third-pass MCAs)	7/82
Engineering prototype runs VMS	8/82
Manufacturing prototype power on	9/82
Two Manufacturing prototypes completed	11/82
All Manufacturing prototypes completed (25)	5/83
Preliminary 102 test done	5/82
Final 102, EMI/RFI & acoustic tests done	12/82
Start field test	12/82
DMT one-third done	2/83
DMT done (CI Base)	4/83
System announceable	5/83
CI Base FVC	5/83
CI Base FRS	7/83
CI Base PA	11/83
FPA FVC	8/83
FPA FRS	10/83
Memory expansion FVC	11/83
Memory expansion FRS	1/84
IDTC Base FVC	3/84
IDTC Base FRS	5/84

Technology

FINCUT (MCA autoplacer) production release	Done
Breadboard conceptual review	Done
Final product conceptual review	Done
MCA socket decision made	Done
Custom register file chip wafer fabricated (second pass)	3/81
Breadboard parts available	3/81
Breadboard mechanical assembly complete	3/81
Breadboard cables available	3/81
Mechanical breadboard power on	3/81
1K and 4K RAMs Qualified	3/81
Modular power system units FVS (Burlington)	6/81
Prototype material available	2/82
Mechanical prototype power on	4/82

MCA-Hudson

Motorola contract signed	Done
Motorola process data to start HL	Done
MCA evaluation complete	12/80
Parts from Motorola production line	1/81
Motorola data base stable	3/81
MCA wafer processing started at Hudson	6/81
Hudson process compatibility verification	12/81
Motorola process qualified	2/82
Hudson process qualified	9/82
All Motorola MCAs at QVL	1/83
7 Hudson MCAs at QVL	3/83
Hudson inventory reaches 1000 parts	5/83

Multi-Signal Layer, Controlled Impedance Program

Board technology center approved	Done
Precision artwork lab completed	Done
Precision artwork lab qualified and tool generation established	1/81
External vendor base qualified	1/81
Product and material specification completed	3/81
Qualification plan for external manufacture	4/81
Transfer buy responsibility to EBB	7/81
Technology center completed	7/82
Qualification plan for internal manufacture	7/82
Volume production (50% of needs)	1/83
Transfer process to volume plant	1/84

Processor and Adapters

First MCA to layout	Done
Machine partitioning complete	Done
First module to layout	Done
Detailed specifications published	1/81
SBIA breadboard power on	4/81
Start console breadboard test	4/81

Peripheral Equipment (FRS dates)

CI780	3/82	TA78	11/82
UDA50	11/81	LCGCR	2/83
HSC50	5/83	DMP11	5/81
RA80	11/81	DZ32	5/81
RA81	8/82	DZ730	11/81
Pinon	8/82		

Software

The software milestones are actually inherent in the hardware milestones listed above and are totally dependent on them. On the prototype, VMS must fully run, so all chips can be released without doubt or qualification. In addition there are these internal VMS milestones, which depend explicitly on the prompt availability of the hardware specifications. The present schedule has minimum contingency to tolerate any additional hardware delays.

Venus VMS functional specification complete	4/81
Venus modifications done	10/81

Diagnostics

Begin diagnostic console debug	5/81
Begin SBIA diagnostic debug	9/81
Begin instruction test debug	9/81
APT and RD support complete	4/82
CPU partial isolation	7/82
CPU full isolation	7/83

Qualification

DVT begins	7/82
Mechanical sensitivity testing completed	8/82
Bus measurements completed	9/82
Functional testing completed	10/82
RAMP verification completed	10/82
Electrical sensitivity testing completed	11/82
Product performance testing completed	1/83
Configuration sensitivity testing completed	1/83
DMT begins	1/83

Manufacturing

Manufacturing product/process strategy	Done
Manufacturing plan completed	6/81
Capital equipment funding	9/81
Start process validation (build QV)	4/82
Implementation plans completed	11/82
Start production	3/83
Dock merge system certification (first packaged system)	2/84

Customer Service

Set Phase 1 BMC goals	Done
Forecast FRS spares	8/82
Start training FRS personnel	8/82
Final RAMP review	1/83
Start formal FS training	6/83

3.5 RISKS AND DEPENDENCIES

Following are some of the more critical risks, whose impact on the Program could be severe. In each case whatever actions can be taken to lessen the risk are indicated, and backup strategies are considered wherever appropriate.

Macrocell Array

With MCAs, the most significant risks are the ECL performance and degradation numbers. We have completed a detailed evaluation of eight test lots from a prototype line, and the results have been encouraging. We will continue simulation efforts in an attempt to determine the worst-case numbers as early as possible, and we will also be evaluating our own components from a production line so we will know how well all AC specifications can be met before the delay analysis of the machine begins in April, 1981. The only possible backup is to accept poorer AC specifications, as 100K ECL logic is not regarded as a viable alternative to MCAs. There is also some risk in adverse impact on the schedule by having Motorola as a sole source, but we are compensating for this by bringing the MCA process into the Hudson facility.

Multi-Signal Layer, Controlled Impedance PC Boards and Backplanes

Boards and backplanes with multiple signal layers are absolutely necessary to provide sufficient component interconnectivity to allow the required component density. Such units are in general use, but no vendor can supply them in the quantity needed or at a suitable price. The risk here is in acquiring the know-how and capacity to manufacture them in-house. The Corporation is now convinced that this is the path to follow, and we will assist in any way we can in developing the necessary capabilities. There is sufficient outside capacity for the first year of Venus shipments, but we must have in-house production from a volume plant by FY85. With this in mind, the MSL/CI Program Team is aggressively pursuing a plan to develop a pilot plant and then get into volume production. In a real sense there can be no backup for

multilayer boards: the system as presently conceived is impossible without them. Backup for the backplane is to have as many layers as possible and take up the slack with twisted pairs; we are developing an interconnect model to determine exactly what is needed and how reliability is affected.

Schedule

The overall schedule is dependent upon the timely meshing of a multitude of individual tasks, and a delay in any one of them is quite likely to have at least some impact on the Program as a whole. Events that would have the greatest adverse impact on the schedule are any failure on Motorola's part to meet their claimed turnaround times (which should be alleviated by having Hudson as a second source), unavailability of 1K or 4K ECL RAMs in the required volume in the timeframe set by our schedule, and any unforeseen contingencies such as illness or accident. The time required for hardware, diagnostic and VMS debug is a function of how well the hardware and microcode are designed; there is some concern in the VMS Group that the allotted time is too tight in light of their experience with the 11/750. However as explained in the preceding section, we have devised mechanisms for tracking the Program to a very fine degree. Hence should any unforeseen event occur, we at least will know exactly what effect it will have on the various parts of the Program, and we can thus take remedial action.

Transfer Cost

Ninety percent of the transfer cost is in materials alone, and in many cases it is the high risk technology itself that is critical to meeting cost objectives: any alternative to MCAs and multilayer boards would be more expensive. Manufacturing is naturally pursuing all avenues for arranging the best possible terms for purchase of the needed materials, and is also investigating any savings that would result from bringing processes in-house rather than purchasing from vendors. Of particular importance is the establishment of volume in-house production of MSL boards, as the capacity is necessary not only to satisfy the need, but also to achieve the cost goal.

MCA Sockets

Use of sockets for MCAs would greatly facilitate replacement in the field, and would result in considerable savings in both manufacture and service. All functional issues have been resolved, and the remaining problem is predictability: the effect on

reliability is not well-known, although it is believed that use of a socket reduces the reliability of an MCA by at worst 25%, principally because of all the extra contact points. For this we are doing life tests to establish reliability numbers.

If the sockets cannot be used, Engineering will hardwire the chip into a plugin package that fits into the same hole layout as the socket and is soldered to the board just like a typical IC.

RAM Sockets

The value in the use of sockets for RAMS is the same as for MCAs, the functional issues have been resolved, and their reliability is already known. Hence sockets will be used for RAMS. However there is an additional problem in that boards without MCAs use half-inch spacing, and the sockets may not fit on the three such boards that do have RAMs. But if the sockets cannot be used in any individual cases, RAMs will simply be soldered to the boards as has been done in the past.

Software Synchronization

The best possible scenario would be for Venus FRS to coincide with a major VMS release, in particular Release 3B. However the current software schedule may not allow VMS to be run on a Venus prototype before Release 3B. The fallback position is therefore to release VMS with partially tested Venus support, and then to provide binary update patches for full support at Venus FRS. On the other hand, any delays in the VMS schedule can only benefit Venus: the longer the delay, the more likely Release 3B will coincide with Venus debug and thus provide full support at that time.

FCC Regulations

Current FCC rulings mandate testing, correction and documentation of all computer products relative to new requirements pertaining to radio frequency emissions resulting from the high speed switching in data processing equipment and electronic pollution of power lines resulting from ineffective or improper line filtering. This will require a major program for LSG. In preparation for the effective date of these regulations, the Technology Group has investigated the matter and issued a report outlining the kind of program required for both new and old products, its cost, and the cost and availability of test equipment and facilities. The report also lists current products that will have to be tested and corrected.

We shall handle the testing for the three cabinets unique to Venus and guarantee that they satisfy the FCC requirements. Moreover we shall work with the 11/780 people concerning the optional second Unibus cabinet and the SBI expansion cabinet. However the principal problem area is the IO equipment, which is under the purview of the Corporate FCC Group led by Joe Smith. Product Management will supply the Group with a list of the options that we wish to support, and our technical people will work with the Engineering managers in the IO area to track possible solutions. Overall responsibility for ensuring that products meet both FCC and acoustic standards lies with the individual functional groups: Distributed and Midrange Systems Development and Storage Systems Development; these groups are now preparing their plans, which should be firm by the end of FY81. It should be noted that at this point the test mechanism is not entirely clear, and there are also legal issues that must be resolved, in particular exactly what must be tested: all possible configurations, typical configurations, maximum configurations.

Dependencies

Besides the major risk areas that could have such an adverse effect, the Program and its various parts are dependent for their fulfillment on the performance of many other people and groups throughout the Corporation. Here are some examples.

The design and layout of MCA chips, pc boards and modular backplanes are heavily dependent on the development of very sophisticated software, which in turn is heavily dependent on computer time and personnel to do the job. In addition to schedule risks, there are also risks in the processes themselves. For example, the representations of some complicated designs may be too cumbersome even to handle.

The peripheral strategy is especially dependent on the definition, funding, scheduling and meshing of projects from many other areas of the Corporation. These range from the Corporate Interconnect Strategy, whose failure could leave Venus without cost-effective peripherals and create havoc for diagnostics, down to the availability of the UDA50.

The diagnostic strategy is based on the T-11 microprocessor being available on time and meeting its specified functionality and speed. Diagnostics is dependent on the VOTE simulator being available

on schedule and having the expected functionality; a failure here could prevent verification of fault coverage.

3.6 TIME TO MARKET, SYSTEM CONFIGURATIONS, AND SUPPORTED PERIPHERALS

As explained in Section 1.1, the strategy is to go to market first with a midrange system built on the CI Base, then go on to other facilities for system expansion and greater capability, as well as the smaller IDTC Base systems. FRS for CI-based systems or the Massbus backup is July, 1983. Figure 3.2 shows the volume shipping rate (in units per month) attainable by Manufacturing at the proposed Engineering schedule.

System Configurations

All three of the base systems defined in Section 1.1 use a double width highboy CPU cabinet and an H9602xx Unibus-console cabinet; the CI Base also requires a lowboy cabinet for the HSC50 and RA81, and the Massbus Base requires an SBI expansion cabinet for the RH780s. Expansion possibilities in these and additional cabinets are as follows.

CPU cabinet

At FRS

- Floating point accelerator (FPA)
- Additional memory to 8 MB
- 1 more Unibus adapter for extra IO
- 1 A bus tap for special equipment
- Serial line unit for remote diagnosis

At FRS + 3 months

- 1 more SBIA to add more devices (all adapters on second SBI must be mounted externally, i.e. in SBI expansion cabinet)

Unibus-console cabinet

Contains 1 BAll-A mounting box for DZ730s, DMP11s, DMR11s and so forth, 1 distribution panel, console load device, and optional stepdown transformer; has space for at least 2 more distribution panels (probable maximum 48 lines, but we are attempting to fit 64)

Additional standard 11/780-type Unibus expansion cabinets as needed

RA81 cabinet (H9642) - holds 3 RA81s

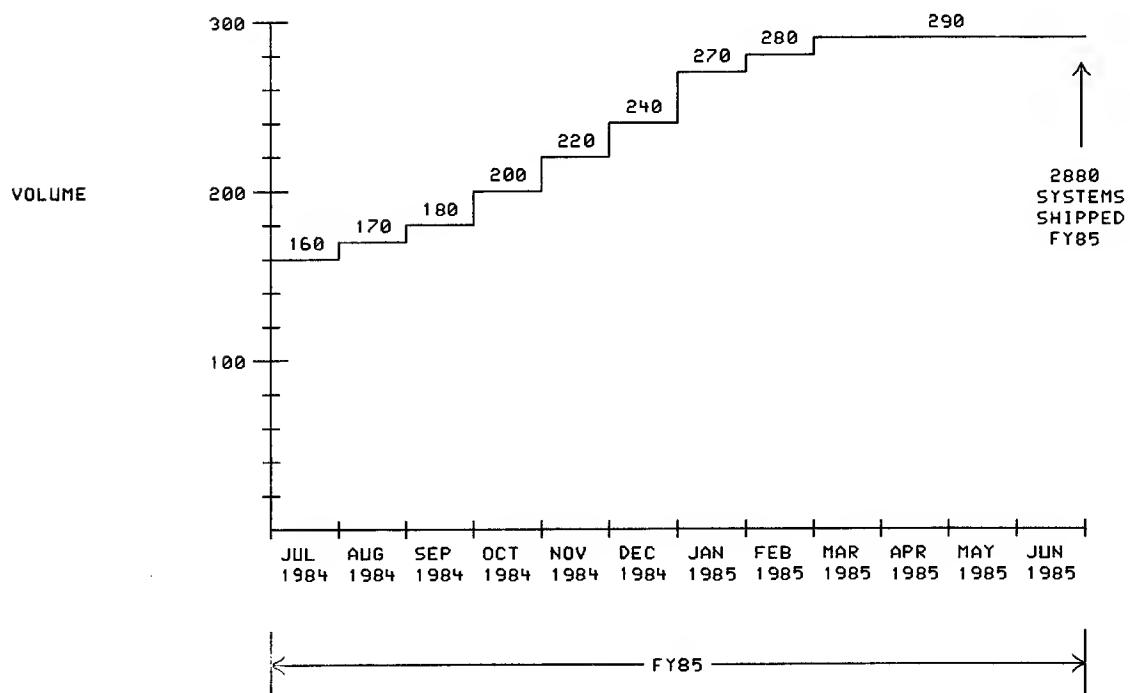
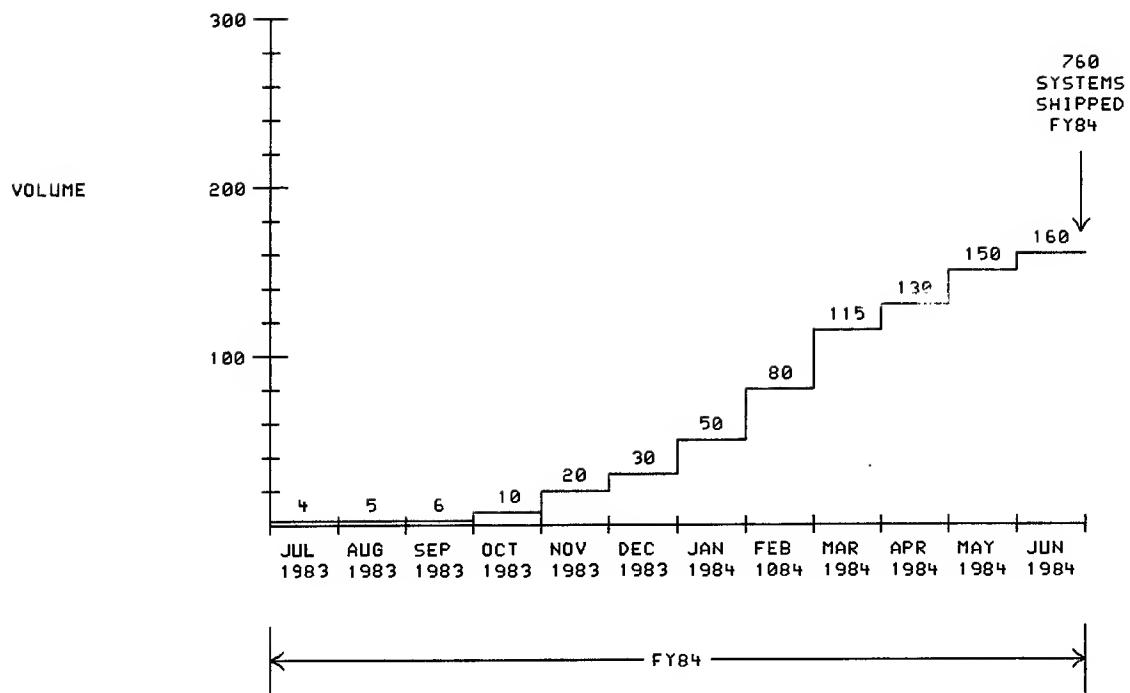


FIGURE 3.2 VENUS SYSTEM SHIPPING RATE

1 or 2 SBI expansion cabinets (at FRS + 3 months)
 Each cabinet can hold 4 adapters connected to the external SBI (a Massbus system requires 1 SBI expansion cabinet); available adapters are DW780, CI780, DR780, RH780

Memory expansion cabinet (at FRS + 6 months)
 Allows for an additional 24 MB of MOS memory with battery backup (and optional stepdown transformer)

Optional Peripheral Equipment

During the first year after FRS, system configurations will be validated with additional options from among those listed here, and no others will be allowed. Options to be supported at FRS are indicated by an asterisk; parentheses indicate options that VMS is not now committed to support, and Product Management will negotiate the status of these with the VMS Group by the end of FY81. Brackets indicate Massbus disks and tapes.

Disks	Tapes	Real Time & Communications	Unit Record
-------	-------	----------------------------	-------------

New products in packaged systems

Pinon*	LCGCR	DZ730*	CR11*
RA81*	TA78*	(KMD11*)	LA120*
[RM05*]	[TS11]	(KMS11*)	(LP07*)
[RP07*]	[TU78*]	PCL-11	LP25
HSC50*		CI780*	LP26*
UDA50			

Products available as add-ins and add-ons

RA80*	TU58*	(COM-IOP/DUP*)	LP11*
[RM80*]		(KDZ*)	
RL02*		DMP11*	
RX02		DMR11*	
		(DN11)	
		DUP11*	
		DZ32*	
		(KMC11-B)	
		DR780	

For field upgrade only (no new ship)

[RP06]	[TU77]	DR11-W
		LPA11*

4 COST

The major cost categories that must be considered are the cost of developing the product, the cost of building it, and the life cycle cost.

4.1 DEVELOPMENT COST

Included in this category are all design expenses, plus the pre-FRS startup expense for Manufacturing, Hudson LSI, and Customer Service. Figures given in the tables on the next three pages are in thousands of dollars.

Hardware and Software Engineering

The hardware budget given on the next page includes material expenses to cover building two breadboard and five Engineering prototype systems, creating all diagnostic programs, releasing forty-five MCA types, nineteen L modules and four backpanels to Manufacturing, and integrating all devices that we have committed to support. Also included are funds for Venus's share of the FCC and MSL Programs, and for Manufacturing process development (assembly and test) as stipulated in the Manufacturing Product/Process Strategy. Following each line item is the name of the responsible manager. Note that the budget does not include funding for building the Manufacturing prototype systems. It is estimated that these will cost approximately \$100K each, and we recommend that X95 manufacturing accounts be opened to capture these costs. Individual user groups must capitalize and depreciate their machines, as the Corporation will charge them off to the various cost centers over five years.

	FY80	FY81	FY82	FY83	FY84	Total
CPU/system design I/E/M boxes, SBIA console, microcode Sas Durvasula	1884	3782	4620	3590	2258	16,134
FPA Sas Durvasula		125	766	622	210	1,723
Memory array design Sultan Zia	242	259	88			589
Memory expansion Sultan Zia		77	266	150	100	593
Current peripheral diagnostic integration Dick Beaven		67	132	145	80	424
Technology						
Circuits/mechanical International regulations	925	1331	1063	390	150	3,859
Register file	34	39	30	35	40	178
MPS	145	125	10			280
		130	10			140
Sultan Zia						
MCA engineering Bill Walton	860	596	500	150		2,105
CAD tools Roy Rezac	115	215	130	72	80	612
Qualification Ron Setera	16	52	115	280	60	523
Prototype/pilot manufacturing & ECO Vic Ku			42	1000	2200	3,242
Release Joe McMullin			100	200	50	350
LSG computer operations Tim Beers	626	1039	1434	1445	1672	6,216
FCC Program (Venus share) Sultan Zia	127	165	36			328
Contingency Vic Ku	93	829	985	1000		2,907
<hr/>						
Total Hardware	4846	8057	10300	9100	7900	\$40,203K
<hr/>						
VMS development Joe Carchidi	15	99	218	120		\$452K
Multi-Signal Layer, Controlled Impedance Program (Venus share) John Belanger						\$1,215K

The amounts given for VMS development (at the bottom of the table) are for design, writing and debugging of code specific to the Venus processor and IO adapters. There are other software expenses related to Venus that are part of central VMS development activities. These include testing and quality assurance that attend a major VMS release, support for Corporate interconnect controllers and devices (available to all VAX systems), documentation of all VMS work (including Venus specific) by the VMS writing group, software distribution, and continuing software support (handling SPRs, developing fixes, DECUS, etc).

Manufacturing Startup

Bob Murphy, John Grose

	FY80	FY81	FY82	FY83	FY84	Total
MR new product startup	127	922	1990	2382	609	6,030
MR capital and tooling				3520	2045	5,565
MO new product startup		41	124			165
MO capital and tooling			139	520		659
Total	127	963	2253	6422	2654	\$12,419K

Hudson Startup

Ken Brabitz

	FY80	FY81	FY82	FY83	FY84	FY85	Total
--	------	------	------	------	------	------	-------

E97 New process							
Manpower	50	600	750	400			1,800
Contract	50	100					150
Material		150	250	200			600
E96 Plant funding							
Manpower	20	66	144	320	340	175	1,065
Material		100	30	210	226	140	706
Masks			48	90	96	60	294
CAD			5	20	22	10	57
Tester hardware			5	10	10	5	30
Qualification		25	50				75
E69 New product startup			40				40
Capital equipment*	150	700	300	800	500		2,450
Burn in (if required)				50	50		100
Total	270	1756	1557	2150	744	890	\$7,367K

*Hipox '80, 3280s '81 & '83, dedicated tester '85.

Customer Service Startup

Mike Robey

	FY80	FY81	FY82	FY83	Total
CS System Engineering					
ME Group	102	149	149	149	548
Spear Group	49	49	49	49	197
Course Development	42	110	159	406	717
Training				290	290
Warranty 30% (87)					
Capital equipment					
MEG lab				100	100
Training lab				100	100
Software ME Group	12	12	12	24	60
Total	205	320	369	1118	\$2,012K
Total Warranty	205	320	369	915	\$1,809K

All of these costs except 70% of Training are born by Warranty, and they therefore show up on the Warranty line of the Project Financial Analysis statement, rather than appearing as Customer Service startup expense.

4.2 MANUFACTURING COST

Here we consider the actual cost of producing an individual system. Figures are in dollars, estimated for the 850th Basic System with the dollar valued at time of production. Also given is a detailed breakdown of the processor based on these vendor costs for shipments in FY84:

MCA	33.20	8-signal-layer module	280
1K RAM	5.50	4-signal-layer module	70
4K RAM	15	Memory register file chip	10
64K MOS chip	9	A bus interface chip	21.60

The above figures are actual vendor quotes, and they can be expected to fluctuate depending on volume, yield, and the general economic situation. As has already been pointed out in Section 3.5, at steady state 92% of the cost of the CPU kernel is materials, and Manufacturing is investigating every possibility of savings, both in purchasing and in developing in-house capabilities. There is currently no firm

Second-source quote (Hudson). The projected MSL cost is expected to decrease considerably in FY85 with in-house production from a volume plant. Note that the following costs are contingent on QBON chip isolation.

Processor

Double width cabinet with power	\$6912
Battery backup	475
I/E box modules (10)	12658
M box modules (3)	3769
Console module (1), LA120	1542
CPU/memory backplanes	2140
IO backplane	749
Memory bus terminator	316
A bus & SBI terminator	100
1 SBI adapter with A bus backplane	1176
Assembly (17.7 hours)	797
CPU test (9.5 hours)	509
<hr/>	
Total	\$31,143

Unibus-console Kernel

H9602-xx cabinet	600
1 BA11-AE mounting box	1470
3 DD11-DK backplanes	354
2 M9202 jumper cards	54
1 M9602 Unibus terminator	39
27 G727 continuity cards	73
1 DZ730 Combo with distribution panel	800
1 RL02 disk	1095
Assembly (4.5 hours)	202
Test (1 hour)	54
<hr/>	
Total	\$4,741

CPU Cluster

Processor	\$31143
1 MB memory	2004
1 DW780xx Unibus adapter	1200
1 Unibus-console kernel	4741
Cables	80
Packaging	150
<hr/>	
Total	\$39,318

CI Base

CPU cluster	\$39318
3 MB additional memory	6012
1 additional DZ730 Combo	695
1 CI780 adapter	2455
1 HSC50 controller	6800
1 RA81 disk	4200
1 KSTI HSC50 subcontroller	600
1 TA78 tape	13100
Total	\$73,180

IDTC Base

CPU cluster	\$39318
1 DW780-UDA50 disk-tape controller	2323
1 LCGCR tape	6000
1 RA81 disk	4200
1 Pinon disk	3000
Total	\$50,641
	\$51,841

Massbus Base

CPU cluster	\$39318
1 MB additional memory	2004
1 additional DZ730 Combo	695
1 H9602-HA SBI expansion cabinet	1828
2 RH780-AA/RH780 Massbus adapter	3783
1 RP07 disk	10700
1 TM78-TU78 tape	13500
Total	\$71,828

Optional Equipment

FPA	3699
1 MB memory array module	2004
Memory repeater module	500
Battery backup for 12 MB	475
Memory expansion cabinet (including backplane)	5925
Pinon disk	3000
LCGCR tape	6000
LP26 line printer	5000
SBI adapter	766
H9602-HA/HB SBI expansion cabinet	1828
Stepdown transformers	
CPU-Unibus cabinets	1300
Memory expansion cabinet	1000

4.3 LIFE CYCLE COST

A fundamental concept underlying the program to create the Venus system is to minimize the total cost of that system over its entire life while at the same time promoting the greatest possible customer satisfaction. With this objective in mind, in August, 1979 Win Hindle formed a Venus System Life Cycle Cost Task Force (Ulf Fagerquist, chairman; Walter Manter and Dave Thorpe, members). The Task Force is studying the total life cost of engineering, manufacturing, installation, warranty and service for Venus systems with a view toward achieving these three objectives:

- a) To develop methods and define parameters for establishing and forecasting levels of customer expectation and satisfaction.
- b) To develop methods for minimizing the total cost and maximizing customer satisfaction for the Venus system, especially methods that do not require additional expenditures.
- c) To develop an integrated financial model as a tool for evaluating the efficacy of the methods developed in b).

The results of the efforts outlined above will of course be general in nature and applicable to any project. For the immediate situation, the key objective of the task force is to set the direction and give proper guidelines for the Venus Program. These guidelines should give visibility to the need for commitment to Corporate-wide processes for the planning, integration, and implementation of the Program across all functions: sales, product lines, services, manufacturing, engineering, etc.

In conjunction with the Task Force, Venus Program leaders have already generated a list of desired alternatives related to life cycle cost. Dates for resolution of these alternatives have been set ranging from 3/80 to 12/81 depending on priorities and needs. Some of the more important alternatives and the dates that they were or are to be resolved are these.

HPP vs conventional mechanical packaging (conventional packaging selected)	3/80
Sockets for RAMs (will be used)	7/80
Sockets for MCAs	12/80
Level of dock merge	6/81
Burn-in of components	6/81
Provide module level specifications	6/81
Improve Field Service reports	12/81

5 SYSTEM REALIZATION STRATEGIES

It is one thing to define a logical system and even to develop it into a physical entity, but then the result must be made a reality - it must be manufactured as a marketable, maintainable product. Following are the strategies for accomplishing this objective.

5.1 MANUFACTURING STRATEGY

Manufacturing involves special considerations for the components. The MCA build process developed by Motorola has been purchased by the Corporation and transferred into our Hudson facility. The MCAs will be tested by both Motorola and LSI Central Incoming Test using the D-LASAR program. Incoming 1K and 4K RAMS will also be tested at LSI. Remaining ECL components are expected to be high reliability parts.

The modules for Venus will be built, tested up to QV level, then integrated in volume with backplanes, power supplies, cabinet, and memory to form a CPU kernel. This kernel is intended to be dock mergeable. It is expected there will be three categories of system: one that is dock mergeable out of the volume area; one requiring some custom configuration; and one requiring full custom configuration. These three categories are referred to as dock-mergeable systems, packaged or common systems, and complex or a la carte systems.

Dock-mergeable systems will be assembled using off-the-shelf options from the stockroom, and it is expected that time to ship from receipt of customer order will be one week. The dock-mergeable stock kept on hand will be determined by marketing forecasts.

Based on marketing forecasts, Manufacturing will build a number of packaged systems of various

types documented and verified by Engineering. Then upon receipt of a customer order, the appropriate system will be taken directly from finished goods stock, with possible add-on of perhaps one or two dock-mergeable custom options. This category will require two weeks for FA&T: one for packaging the system before hand, and a second to ship following customer order.

Complex systems will be assembled individually from finished goods stock upon receipt of customer order. The time required will be four weeks.

It is expected that these three types of process will carry FA&T costs of \$3,000, \$7,000 and \$14,000 respectively (based on cycle times of one, two and four weeks).

Marlboro has been selected as the plant for both volume and FA&T. Charlie Bradshaw, the site Manufacturing Manager, will be addressing all issues relating to Venus manufacturing.

5.2 DOCK MERGE

Both Engineering and Manufacturing feel that considerable savings in FA&T costs can result from having as much of the system as possible be dock mergeable. Manufacturing has published an analysis showing that an aggressive dock-merge program would result in a total cost avoidance to the Corporation of \$99 million over the life of the Venus product. Each configuration will however be investigated to determine whether dock merge is worthwhile in light of any offsetting customer dissatisfaction, higher installation or warranty costs, etc. For an option to be dock mergeable means it is expected to perform at an IQ level of 90% when incorporated into a system that is dock mergeable. To this end we will do the following.

We will make certain all of the equipment we design is dock mergeable. At FRS the basic dock-mergeable CPU kernel produced by Manufacturing for all systems will be the CPU cluster plus 1 MB of memory. This is appropriate for all systems that will be delivered in the early stages, and Manufacturing is prepared to create a smaller kernel should that become necessary. For final testing each kernel will include an FPA, but it will then be removed, and both it and the kernel will be stocked separately as dock-mergeable options.

We have already determined which peripherals will be supported at FRS and which at FRS + 12 months; of these, which are listed at the end of Section 3.6, all will be dock mergeable. Dock merge is now a Corporate goal for all new products, and we will attempt to make maximum use of such new products. Furthermore, wherever feasible, we will do our best to influence the individual product groups to qualify whatever of their current products that we may decide to support in the future.

Manufacturing is currently conducting an analysis of its procedures with a view toward making whatever modifications are necessary for dock merge. Areas particularly affected are the process, material and capitalization plans.

For the success of dock merge, Manufacturing is dependent on Qualification Engineering and Customer Service. Qualifying a product for dock merge relies principally on the various maturity tests and the many qualification tests outlined in the next section. Jointly with Customer Service, Manufacturing will conduct sample audits of specific products during installation, and will establish a formal feedback procedure for information on whether dock-merged systems are living up to their expectations in the field.

We will define a reasonable number of packaged configurations of various sizes to allow the greatest degree of dock mergeability, consistent with reasonable cost and flexibility in tailoring individual systems to customers' needs. For this purpose we will maintain strict revision control in order to know the status of all elements at all times.

If this endeavor is successful, peripheral controllers and extra memory array boards can be tested on a single CPU kernel kept right on the test floor, rather than needing to be tested with the particular options that will appear in a final customer system. This will of course eliminate a great deal of handling. The various options will be kept on the shelf, and final assembly of a customer system can be accomplished by merging a CPU kernel with the FPA, extra memory, peripherals, and the software package.

Even for systems including devices that are not dock mergeable, there is no need to assemble the individual customer system for final test. Instead the FA&T area will have several option test stations available on the test floor. Then for a special

system, the nondock-mergeable items can be tested on these stations.

5.3 QUALIFICATION PLAN

This plan defines the steps required to ensure that Venus is a qualified product. To be qualified the system must:

Adhere to Engineering functional and performance specifications.

Achieve product cost goals.

Achieve MTBF and MTTR goals.

Undergo a formal DMT and PMT to prove that the system lives up to defined RAMP goals.

Qualification approaches the product from a system level to prove that this system can perform the functions previous systems have and to build confidence in areas of risk. A risk area is any as yet unproven functionality or technology. To reduce exposure to risk, vulnerable areas will be identified and provisions made to test them. A heavy emphasis will be placed on ensuring that all system RAMP features are achieved.

Testing procedures fall into two general categories, system qualification and product performance. The schedule for these tasks is determined principally by the schedule for the Program as a whole.

System Qualification Tests

These are tests to confirm that Venus lives up to its design specifications in the areas of sensitivity to various conditions, conformance to DEC standards, functionality of various procedures and features, and characteristics of the overall system. Tests are performed in the following categories.

Sensitivity of the system to margins in voltage, temperature, humidity, and clock rate.

Conformance of electrical characteristics to international regulations, DEC Standards 102.7 (particularly EMI/RFI), 122 and 123, and FCC requirements.

Conformance of mechanical characteristics to international regulations and DEC Standards 102 and 119.

Sensitivity of environmental and power sensors.

Verification of performance under voltage margins.

Verification of all power up/down sequences.

Verification of initialization and configuration routines utilized to prepare the system for bootstrapping.

Verification of the bootstrap program via all load paths.

Verification of software support of the hardware (i.e. functionality of the operating system).

Verification of the orderly shutdown of the operating system.

Verification of the diagnostics used in the manufacture of the system, other than those associated with special test equipment, which will be verified by Manufacturing.

Verification that the diagnostics provide the tools necessary to maintain the system.

Verification of system RAMP features.

Determination of the sensitivity of the hardware and software to variations in system configuration.

Determination of the reliability of the hardware and software, especially RAMP features, through the design maturity test (to measure MTBF) and a software load test.

Product Performance Tests

Hardware and software benchmarks will be run to measure the performance of the system against existing Digital computers and those of the competition.

5.4 DIAGNOSTIC STRATEGY

The goals of the Venus diagnostic strategy are to detect in excess of 99% of all solid faults in the CPU cluster and to isolate the detected solid faults to

the failing module 95% of the time. For components outside the CPU cluster, the corresponding goals are 95% detection and 90% isolation to the module level. A further goal is to do the isolation in 2 minutes or less, exclusive of any time needed to handle media.

To achieve these goals, the Venus diagnostics are organized into two types of test, Functional Fault Detection, which are in Venus microcode, and End State Verification, which are driven by the T-11. Both have several isolation modes and make use of "tick files", which are the recorded behavior of a known good machine after each clock tick.

Further goals of Diagnostics are to detect and report the symptoms of intermittent faults as accurately as the hardware permits, and to supply a single console software package that supports both standalone diagnosis of a failing system and normal operation under VMS.

To support the manufacture of Venus, the repair diagnostics will be written to be completely self-contained and independent. Thus the strategy is to run only those tests that involve logic on the board under test at the QV station.

To support Engineering design verification of the CPU prior to layout, register transfer level models of the I, E, M and F boxes with input test vectors to exercise them will be provided. We shall also provide functional tests to aid in breadboard and prototype checkout, to verify that the CPU and its microcode do indeed implement the VAX architecture, and to assist in verifying ECOs, screening modules, and system acceptance testing.

The quality of the diagnostics will be checked by using the VOTE simulator to run them on a model similar to that used by SAGE2. This actually enables the diagnostics to do fault insertion in the simulation, to verify the level of fault detection and chip isolation in an automated fashion.

Since the bus adapters are to be implemented in TTL rather than with MCAs, the diagnostic strategy is to isolate to a failing module using microcoded functional diagnostics. Existing peripheral diagnostics will be used where possible, and where none exist Large System Diagnostics intends to act as the key interface in getting the appropriate groups in Maynard or Tewksbury to do the job.

Dependencies

To carry out this strategy, Large System Diagnostics is dependent on other groups as follows.

Hardware Engineering to keep their schedules and provide the diagnostic logic necessary to implement our strategy.

Hardware Engineering and Marlboro Operations to make available 2-4 hours per day on the Engineering breadboard and 6 hours per day on the Software breadboard from power on through prototype power on for console and diagnostic debug.

Hardware Engineering to make available 8 hours per day on a prototype machine for diagnostic debug.

LSG Management for funding in a timely enough manner to allow hiring and training of the necessary new personnel.

Marlboro Operations for providing adequate DECsystem-20 timesharing facilities and a sufficient amount of time (to be specified) on a VAX 11/780 with 8 MB of memory for doing VOTE simulations.

VAX DS Group for providing the Venus version of the Diagnostic Supervisor in a timely fashion.

VMS Group for supplying drivers to incorporate into the DS and for timely availability of a version of VMS that fully supports Venus and its peripherals.

5.5 CUSTOMER SERVICE PLAN

The Customer Service Plan defines a set of installation and maintenance commitments based on an established maintenance philosophy. These commitments are listed in Section 1.1; the philosophy supporting them is given here, followed by a short discussion of the functionality of some of the major RAMP features.

Maintenance Philosophy

The general approach to corrective maintenance is that symptoms of malfunctions will be investigated first at the operating system level (it is in fact the first level maintenance tool). Information pertaining to the malfunction should guide maintenance activity to a particular area of the system for further diagnosis.

This general or specific area will then be diagnosed, in as few modes as possible, until the malfunction is resolved to a field replaceable unit (FRU).

Materials

We will stock at least one full set of modules at the branch level, and this will be backed up by safety stock at the Corporate level. With RAM and MCA replacement a reality, the branch will also stock individual ICs and will be less likely to increase the number of modules as the population expands. Our goal is to ensure 99% reliable operation of spare modules. To reach this goal we will test them in-house in either standalone mode (new modules) or under the operating system (repaired/reworked modules).

Manpower

We plan a high degree of specialization at the branch level and expect to take advantage of 11/780 personnel during the startup phase. General system troubleshooters will be at the DDC and in the Branch, District, Regional, and Corporate Support Groups.

Training

In conjunction with manpower recruitment we will establish both unit and subsystem specialist courses as well as system level troubleshooting courses. The expected length of the CPU-cluster specialist course is five weeks, and the system level or support course will take about ten weeks.

Technical Documentation

Technical Documentation will support both the educational effort and the field maintenance effort. We expect to provide documentation in the form of print sets, maintenance manuals and procedures, site prep/installation/acceptance guides, microcode flow charts, IPBs, and theory of operation.

Remote Diagnosis

The current strategy for utilizing RD is to have the customer call the Digital Diagnostic Center in Colorado for initial diagnosis of a malfunction. The DDC will utilize a host computer to carry out the diagnosis of many systems in parallel. Once the DDC has determined that a failure exists and has isolated it, the local branch specialist will be dispatched with the proper set of FRUs to correct the problem. Software Service plans to utilize remote connection capabilities almost exclusively.

Major RAMP Design FunctionalitySystem Error Detection, Logging, and Recovery

The operating system software, which includes certain system-critical console software processes (and possibly system firmware) will take a predefined course of action upon detection of any hardware or software error. This specifically includes recording various amounts of data, which will be utilized for immediate recovery processes. This data will also be analyzed by Customer Service with the new System Error Analyzer, SPEAR, for maintenance activities.

Hardware Error Detection Logic

All system components are being designed, where possible, to include hardware error detection or correction (not necessarily ECC) on all buses, internal data paths, RAMs and MOS memory. The current assumption is that error checking will cover at least 40% of the expected hardware failures. This error checking, along with the gathering of pertinent hardware status by firmware, the console or the operating system, should enable the system to isolate intermittent errors to one or two modules in the CPU cluster. For other parts of the system, isolation should be at least to the unit (controller or peripheral).

Remote Diagnostic Link

The console design, in conjunction with the console software and operating system software, will support a remote diagnostic link. This will be used for both standalone diagnosis and system level control and monitoring.

Environmental Monitoring

The console will monitor the thermodynamic and power systems in such a manner as to provide operator notification of most typical temperature problems. The power system will be monitored to the extent needed to isolate the cause of a malfunction to the power control or individual regulator.

6 FUNCTIONAL CHARACTERISTICS

Having explained what the system is and how we intend to develop it, let us now consider in somewhat greater detail just how it is expected to work.

6.1 CPU CLUSTER

An introduction to the functionality of the central part of the system has already been given in Section 2.1. The internal logic of the processor - ALU, data paths, RAMs, buses - is organized in terms of longwords of thirty-two bits, with transfers to and from memory in blocks of sixteen bytes (four longwords). From the program point of view the fundamental orientation of the machine is toward bytes, but with complete capacity for handling words, longwords, etc with arbitrary byte boundaries, although operation is most efficient when operands are kept aligned with the memory byte blocks.

I, E and F Boxes

The 8K x 84 writable control store is in the E box; each of the other boxes has a small control store with special microcode for its own operations.

The I box continuously prefetches the instruction byte stream and stores it ahead in an 8-byte instruction buffer, eliminating most of the performance penalty imposed by instructions crossing physical memory boundaries. Op codes and specifiers are decoded from the instruction buffer. For initiating instruction execution in the E box, the op code selects a location in a dispatch RAM; there are two of these, for native and compatibility modes (i.e. for the VAX and PDP-11 instruction sets respectively). The I box also contains an ALU and a copy of the general purpose registers (GPRs) for calculating addresses.

The E box contains the major part of the microcode, namely that which handles the execution of the instructions. The unit is based on a 32-bit binary/BCD ALU and a 64-bit left shifter and data packer/unpacker, both of which receive input from a pair of multiplexers. Available inputs are the operand bus, a pair of 256 x 32 scratch pads (each of which contains a copy of the GPRs), the Q register associated with the ALU for multiplication and division, and the W register, which receives input from either the ALU or the shifter. Besides instruction execution, the E box also handles machine initialization, interrupt processing, memory management, fault and error processing, and console support.

The F box is a dedicated microengine for performing very fast execution of a subset of the floating point instructions. It contains an ultra high speed multiply data path, add-subtract data path, and normalization unit. Hardwarewise this includes two copies of the GPRs, a 64-bit shifter, and a double, pipelined data path for handling the less significant and more significant fractions. This path creates and sums partial products of thirty-two bits eight times per clock tick (33 ns), yielding a single precision result in 200 ns.

Memory Subsystem

This subsystem includes the storage array boards (1 MB per board) and the M box, which contains not only all of the control, transfer and error logic for the storage array, but the data cache as well. The basic storage unit is a block of four 39-bit words, each containing four data bytes and a 7-bit error correction code.

The cache is 2-way set associative utilizing a writeback storage algorithm. The block size is sixteen bytes (each with a parity bit) and the total storage capacity is 16K bytes. Associated with each of the 256 blocks are valid and written bits and a tag that identifies the block. Replacement is on the basis of the least recently used entry, and write allocation is by block to simplify IO reading. Special logic is included for byte write, significantly decreasing storage access requirements. When memory data containing a corrected error is placed in the cache, the written bit is turned on to force eventual rewrite of the memory location, thus reducing the probability of a double error.

The control logic includes a translation buffer or page table for translating virtual into physical addresses. The buffer contains 512 entries each for system space and process space.

Console

The console is actually a microprocessor-based subsystem. It monitors environmental and power supply conditions, serves as the VMS operating system terminal, and provides an assortment of diagnostic functions. Via the I bus the microprocessor initializes and bootstraps the system and implements such functions as examine, deposit, start, halt. The serial diagnostic bus is used to initialize various diagnostic operations and monitor backplane signals and other error and diagnostic conditions. Besides the microprocessor and its associated logic, the hardware includes a PROM for the bootstrap code, a control RAM for the console microcode; controllers for the console load device, console terminal, and remote diagnostic link; and the environmental and power monitors.

6.2 PERIPHERALS

Hardware in the area of peripherals includes many already existing products or committed projects whose functionality is already known or defined elsewhere. We therefore concentrate here on the adapters, controllers and devices whose development is specifically for Venus.

A Bus Adapters

The SBI adapter, implemented with TTL MSI logic, interfaces the A bus to the SBI. It provides all of the SBI functions necessary for handling CI, Unibus, Massbus, and other devices available on the 11/780 SBI, except that it is not intended to serve as a direct connection to any other processor. The hardware provides protocol, timing and termination for the SBI, and includes registers and data assembly facilities for transfers in both directions.

Mass Storage

In addition to certain enhancements to already available controllers, the Venus Program expects to make use of the following two controllers whose functional character has been determined. The mass

storage devices for these controllers are covered in Section 2.3.

HSC50 ~ This intelligent controller contains six subcontrollers for disk and tape systems. It is based on a microprocessor that optimizes the data buffering, continued IO overlapped with error handling, and overlapped or ordered seeks. It also provides many RAMP features, such as comprehensive error detection and recovery, bad-block revectoring, online volume backup, self-contained diagnostics, online drive repair, a remote diagnostic connection, and support for logical redundancies.

UDA50 ~ This intelligent controller contains four subcontrollers just for disks. It has the basic features of the HSC50 for low cost, medium performance systems.

Communication and Unit Record Equipment

The Venus Program recommends a number of projects in this area: an NI Port Adapter, interfaces for synchronous lines, asynchronous lines, and unit record equipment, and a terminal.

NI Port Adapter ~ NI adapter implementing VAX port architecture.

Synchronous Line Interface ~ A suitable NI interface would be capable of handling two lines, each of up to 38 kb (X25) or 19.2 kb (DDCMP), or one X25 line up to 56 kb. Protocols downline loadable.

Asynchronous Line Interface ~ No async NI option specification exists as yet, but recommended attributes for one are DMA capability outbound with adequate silo, data rates to 9.6 kb, full modem control, and split baud rate, preferably on a per-line basis.

Unit Record Interface ~ No specification yet exists for an NI unit record option, but characteristics recommended are character formating and tab expansion capabilities, and printer support to 1200+ lines per minute. Card reader support would also be useful. Controller should be integral to line printer or card reader.

NI Terminal ~ Like VT100 with transfer cost no more than \$500.

6.3 SOFTWARE

At present the only system software functionalities identified as needed for Venus are in the executive layer. These include support of the processor, IO adapters, RAMP features and console; bootstrapping and initialization; and error handling procedures including error logging. The VMS project for Venus will do the following.

Modify the VMS software bootstrap procedure to configure Venus IO adapters, IO address space, and memory address space.

Implement loadable Venus-specific code for handling run-time processor dependencies. This code will include the capability for handling machine check errors, for saving and restoring nonarchitectural processor registers on power fail/recovery and bugcheck, and for handling errors from IO adapters.

Modify SYSGEN for Venus specific testing of IO configuration.

Implement a driver for the console load device.

Enhance the error reporting mechanism to allow logging of errors detected by the console.

Support booting from any new devices required by Venus, such as a system disk.

Support two SBIs.

Integrate all Venus changes into the VMS master sources and system build procedure.

In general any enhancements that would be visible to the user would take place in the bundled layer. Recommended features (which would then be common to all VAX systems) are:

Support for new peripherals and of an IBM channel interface for PCM devices.

Enhancements to the VMS scheduler, including perhaps scheduler classes and ability to specify the minimum and maximum time allocated.

Enhancements to the command terminal interface, user help facility, and operator interface.

Layered Products

The VMS Group does not directly handle unbundled products, but VMS management is responsible for setting system goals and providing the overall system plan. Hence the Group would be involved in any projects recommended for Venus, which might include enhancements to the Fortran compiler for better optimization, creation of a DBMS system, additional languages such as ADA, and support of additional communication ports to other vendors. Of course most of the required layered products are already under development, as indicated in Appendix C. Basically these products fall into five categories.

Language Processors. By VMS Release 2 the following languages had been implemented with native mode compiler and execution: Fortran, Cobol, Basic, PL/I, Pascal, Bliss. Also VAX11 DSM is available as a native mode interpreter, and Coral-66 as a compatibility mode compiler with native mode execution. New languages in the planning or development stages are ADA, APL in a common implementation with the DECsystem-10/20, Dibol which is important in the COEM market, and C, the system implementation language developed by Bell Labs for UNIX. In most cases, development is implemented to the latest ANSI or "de facto" standard plus enhancements that improve competitiveness. New releases will generally use new features of the VMS release and will therefore be synchronized to it, but this is not an actual goal. The VAX11 Basic will be continually emphasized as the vehicle for RSTS/E users to achieve compatibility with Basic-Plus and Basic-Plus-2.

Communication Products. During FY81, DECnet-VAX will become a full Phase 3 implementation. In addition the packet switching network interface (X.25 etc) will be incorporated into DECnet so that its presence and use are transparent to the user. The only significant future development for communication with the products of other vendors is expected to be on an interface to SNA.

Data Management. RMS is bundled with VMS and will be continually enhanced and integrated with a common data dictionary (CDD) and DECnet. A Codasyl-compliant DBMS will be released prior to Venus FRS and development is due to start on a relational DBMS during FY82. Datatrieve will continue to be the standard query language and report writer, and it will be enhanced prior to Venus FRS to interface to VAX11 DBMS, DECnet and the CDD.

General Utilities. Most emphasis here will be on programmer productivity and ease of use. Forms management will be greatly enhanced and its performance improved by merging FMS into the common application terminal subsystem program (CATS/TSS). A second phase, the terminal management subsystem (TMS) will provide a higher level of terminal transparency and performance through the use of 11/23 video terminal concentrators. Transaction processing is also expected to be greatly enhanced by CATS and additional layered software. The Hydra environment, in which Venus can be used, is a whole set of layered products.

Applications. Individual Product Lines will continue to offer applications and tools specific to their marketing efforts, particularly ESG, ECS and GA, but it is expected that most application software will be provided by joint marketing arrangements with third parties.

6.4 TECHNOLOGY

Venus is taking advantage of technological innovation in a number of areas: LSI logic, multilayer controlled-impedance backplanes and circuit boards, modular power supplies, mechanical packaging and environmental sensing. Following is a discussion of the functionality of some of these innovations.

Macrocell Arrays

Until now the semiconductor industry has used three approaches to meet the demand for LSI digital circuits: standard off-the-shelf circuit families, custom circuits, and gate arrays. Standard circuits are very economical, but are not sufficient for the complex, specialized functions that Venus requires. Custom circuits, on the other hand, are quite expensive and regularly have a turnaround time of a year or two. Gate arrays have a shorter turnaround time since the basic array can be fabricated up to metalization, but the interconnecting metal makes the chip larger and increases propagation delays. To provide greater flexibility in circuit design and development, Motorola has created the macrocell approach to custom LSI. This approach circumvents the cost and time factor of custom circuits and reduces the deficiencies of conventional gate arrays.

The macrocell array is actually an extension of the gate array concept. Instead of gates, however, each cell in the array contains a number of unconnected transistors and resistors. Stored in a

computer are the specifications for creating interconnecting patterns that can transform the unconnected transistors and resistors within each cell into SSI/MSI logic functions or "macros". These macros take the form of standard logic elements such as dual type D flip-flops, dual full adders, quad latches, and many other predefined functions. All of these are ECL structures for optimized performance.

The cell library contains 85 macros: 54 for major cells, 14 for interface cells, and 17 for output cells. A single array can contain 106 cells: 48 major, 32 interface, and 26 output. If full adders and latches are used in all the cells this means a single MCA may contain up to 1192 equivalent gates; if flip-flops and latches are used in all the cells there may be up to 904. Typical power dissipation is 4 watts, 4.4 mW per equivalent gate. Contributing to the high performance of the system as a whole is the extremely low propagation delay in major and interface cells: 1.3-1.8 ns maximum, compared to 3.5-6 for the 10K ECL used in the KL10. Also of considerable importance is the high density: 100 gate equivalents per square inch, compared to 20-30 for MSI. This reduces interconnect delays, thus further enhancing performance, and it also lowers packaging costs as well.

Besides MCAs, technologies considered were the Siemens gate arrays and the Fairchild 100K MSI logic family. The criteria for decision were naturally cost, performance, time to market, multiple sources, and the expected future of the technology. All three technologies meet the performance requirement (3.5 times 11/780), but in all other categories the decision is between MCA and 100K. The 100K will be available slightly earlier but at 20% greater cost; moreover MSI is definitely not a technology of the future. One problem is second source, but the contract negotiated with Motorola gives Digital the technology/process transfer and license to enable internal second sourcing.

Circuits

A goal of circuit design is the elimination of all wires on boards and at least 90% of them on back-planes. This will be accomplished by using 16-layer backplanes and 8-layer pc boards (four signal layers). A further refinement is much greater restriction on the variation in impedance per unit length of etch. We will also increase the number of signal pins available on the edge connectors by adding supplementary connectors to the board solely for power distribution.

Power System

The Venus power system is configured from various elements of the modular power system (MPS) presently under development by the Power Supply Engineering Group. The MPS consists of an ac-to-dc input module from which several dc-to-dc output modules are powered. The input module rectifies the ac line voltage into a raw 300 Vdc bus. Each output regulator uses constant frequency pulse-width modulation at 50 kHz to convert the 300 Vdc to the desired output voltage. Input modules being developed are 2500 watt three phase and 1200 watt single phase, of which only the former is presently being considered for use in Venus. Regulators being developed are 5 volts at 200 and 85 amperes and 2 volts at 85 amperes. Outputs can be paralleled allowing tailoring to specific requirements without excessive unused capacity. Besides attaining efficiencies greater than 70%, use of a 50 kHz switching frequency results in even smaller magnetic components than those obtained at the "standard" 20 kHz switching frequency.

The units can maintain proper operation for up to 20 ms of power loss and down to 180 Vac rms line-to-line input voltage. They are mounted side by side above the Venus logic to be cooled by air exiting from the logic; the configuration is one input module at each end, with seven output regulators between them (a similar arrangement is used in the memory expansion cabinet). Powerup and powerdown sequencing is handled by the console. The system provides extensive interface signals for remote diagnosis; e.g. each regulator can be turned on and off by a diagnostic instruction, and monitoring of a module ok signal determines whether the unit is within operating range. The system also has a standby mode in which only the memory array and refresh logic are on, so memory can be preserved while other parts of the machine are being serviced.

Mounted among the regulators in the power supply rack is a monitor board that consists of a pc board and panel, and contains the electronics, fuses and indicators (mostly leds) for both the power system and the environmental sensors. Dc logic voltages will be monitored either by analog comparators or A-D converters; in either case the information is supplied to the console for appropriate action.

A 48-volt battery with builtin charger will back up the memory storage array boards for 10 minutes. The battery output is converted to 300 Vdc and all other regulators are disabled except the +5V unit used for main memory. Also available is battery backup for 100

hours for the time-of-year clock, which is normally powered from an auxiliary +12V output of the input module. In a long power outage, the user may prefer not to run the battery all the way down, as the recharge time is considerable and there may then be no backup for a subsequent minor outage. Thus the system includes a feature that allows the user to select a shorter backup period; this means the system can ride out a number of short power failures at a cost of going down or even losing the time of year during a long one.

Battery Memory	Time Toy Clock	Recharge Time
10 minutes	100 hours	14-16 hours
10 minutes	10 minutes	4 hours
5 minutes	5 minutes	2 hours
1 minute	1 minute	20 minutes

Mechanical Packaging

New packaging techniques appear at every level of the system, from the cabinet down to the mounting of components on the boards. A large part of the effort expended in dealing with mechanical and environmental issues is in setting standards and then getting vendors to provide equipment that meets those standards, but various innovations have also been developed by Digital's own engineers.

It is expected that almost all Venus systems will be installed in computer rooms with a raised floor, under which pass not only the cables and other utilities, but also conditioned air. The overall packaging scheme takes advantage of this physical environment by ducting the bottom of the cabinet into this source of cooling air, which passes under pressure up through the card cage. (At installations without a raised floor, air will be drawn in through louvers at the bottom front of the cabinet.) Air then flows by the power supplies, and three blowers mounted at the top of the cabinet expel it at the rear through mufflers that keep the noise level well within the acoustic limit of 60 dbA. The 10K ECL and other standard circuit components are cooled directly by the air flowing between the boards. MCAs however dissipate too much heat to be cooled in this manner. Hence mounted on the MCAs are special sinks that dissipate heat into the passing air with much greater efficiency.

Modules are the new L type that are the same height as hex boards but have three sets of press pins. These pins are not only denser than the fingers on the old boards, but they also protrude through the backplane so cable connectors can be mounted directly on them, eliminating the space wasted by having slots just for connectors. Between the sets of pins are pads that connect to aluminum ground bars that are an integral part of the card cage (the entire frame thus serves as logic ground). Similar bars at the top and bottom of the frame carry dc voltages that are picked up by pads at the top and bottom of the cards. Although cables still connect the power supplies to the card area, use of these voltage and ground bars combined with placing the power supplies above the card cage eliminates a great deal of the power cabling that made the inside of older machines look so cluttered. Some of the module pins must still be used for power and grounds, but the new arrangement leaves 230 pins available for signals on each module compared to about 168 on the old hex boards.

Another packaging innovation is the use of sockets for RAMs and MCAs. Sockets increase base system cost, but they will decrease maintenance cost, yielding a net, discounted life cycle cost saving, because of the convenience with which individual chips can be removed and replaced.

Environmental Monitoring

Located throughout the cabinet are devices for sensing various environmental conditions. The electronics and indicators associated with these devices are on the monitor board mounted in the power supply rack (there is no monitor board in the memory expansion cabinet). In most cases, conditions are reported to the console for appropriate action.

The principal environmental concern is overheating in the logic, as the junction temperature in the MCAs directly affects their failure rate, which doubles with every rise of 20 degrees C. At a junction temperature of 88 degrees C, the failure rate is 360 per billion hours. Measurements show that incoming air with an ambient temperature of about 24 degrees produces junction temperatures in the range of 100 to 107 degrees. To guard against overheating, linear thermistor networks are used to monitor the ambient temperature of the incoming air and the temperature gradient across the card cage. A single network mounted below the cage triggers a warning if incoming air reaches 32 degrees, and it causes a system shutdown at 42 degrees. Any of three networks above

the logic reports a 10-degree rise in temperature across the cage and results in a system shutdown on a 20-degree differential.

Other environmental features include devices for detecting an overheated regulator or failed blower. Overheating of a regulator, whether caused by faulty operation or too high an ambient temperature, closes a thermal switch that shuts down the main power control. Unless accompanied by a temperature problem, failure of a blower is not serious enough to warrant automatic corrective action, but it is still desirable to be aware of such an event; blower failures are detected by Hall effect switches mounted in the blower housings.

Appendix A

VENUS SYSTEM RELATED PLANS AND SPECIFICATIONS

Top Level Plans and Specifications

Date is when Phase 1 plan was available.

	Revision	Responsible Person	Date
Business Plan (Phase 0)		Carl Gibson	11/79
Business Plan (Phase 1)		Carl Gibson	12/80
Product Requirements (Phase 0)		Carl Gibson	11/79
Development or Project Plans			
System	2	Vic Ku	12/80
Technology		Sultan Zia	12/80
MCA Engineering	B	Bill Walton	3/80
CPU System Engineering	3	Sas Durvasula	12/80
IO	4	John Bloem	12/80
VAX/VMS Venus	3.0	Chuck Samuelson	12/80
Product Process Strategy	1	John Grose Graham Swift Larry Cornell	12/80
Customer Service	2	Mike Robey	12/80
Diagnostic Engineering	3	Dale Cook Jeff Barry	12/80
LSI Business Plan		Ken Brabitz	12/80
Product Qualification Plan		Ron Setera	12/80
Product Description	2	Sas Durvasula Tryg Fossum	12/80
Technical Writing Plan	C	Ed McFaden	12/80

Detailed Specifications

	Revision	Responsible Person	Date
I Box Specification		Tom Knight	11/80
E Box Specification		Bob Elkind	12/80
M Box Specification		Bill Bruckert	12/80
F Box Specification		Tryg Fossum	12/80
System Clock Specification		Paul Guglielmi	5/80
IO Adapter Bus (A Bus) Specification		Jim Lacy	11/79
SBI Adapter Specification	.2	Barry Flahive	11/80
A Bus Tester Overview and SBIA Debug Plan		Barry Flahive	6/80
Console Specification	4	Ed Anton	12/80
Diagnostic Bus Specification		Ed Anton	5/80
Memory Array Specification		Barbara Altman	10/79
Main Memory Specification	2	Al Dellicicchi	12/79
Expansion Memory Implementation		Bruce Weaver	12/80
1K RAM Purchase Specification		John Kelly	1/81
4K RAM Purchase Specification		John Kelly	1/81
RC Terminator Network Specification		John Kelly	1/81
Microcode Specification		Tryg Fossum	12/80
Error Recovery Specification		Frank Robbins Tryg Fossum	12/80
Bootstrap Specification	3	Eileen Samberg	7/80
CSM Design Specification	0.0	John Derosa	10/80
Environmental Monitoring Module Specification		Derrick Chin	5/80

Appendix B

BREADBOARD/PROTOTYPE CAPITAL EQUIPMENT

Engineering breadboards and prototypes will first use current interconnects plus the CI. Other new interconnects and new peripherals will be tested on the prototypes 9-12 months later.

All breadboards and prototypes include a Venus CPU; the A bus tester will use an 11/780. All prototypes will use an RX02-BA as the console load device and an LA120 as the console terminal.

Engineering Breadboards

	A Bus Tester	Breadboard 1	Breadboard 2
Date needed	4/81	6/81	6/81
Main user	IO*	Engineering/ Diagnostics	Software
Memory array SBIA	1	2 MB 1	1 MB 2
SBI bus adapters	1 DW780-AA 2 RH780-AA 2 DR780-AA	1 DW780-AA 2 RH780-AA 2 DR780-AA	1 DW780-AA 4 RH780-AA** 1 IPA780
Unibus equipment	1 DZ11-E 2 VT100-AA 2 DMC11-AL 1 DMC11-MA 1 DMC11-MD 1 H9602-DA 1 H9602-HA	1 DZ11-E 4 VT100-AA 1 LP11-WA 1 DMC11-AL 1 DMC11-MD 1 H9602-DA 1 H9602-HA	1 DZ11-E 4 VT100-AA 1 LP11-WA 1 DMC11-AL 1 DMC11-MD 1 H9602-DA 1 H9602-HA
Massbus equipment	1 RP06-BA 1 RP07-AA 1 TU77-AB	1 RP06-BA 1 TU77-AB	2 RP06-BA** 1 TU77-AB
Console		1 LA34-AA 1 RX02-BA	1 LA34-AA 1 RX02-BA

* For test period only, then most equipment will be moved to Prototype 5.

** 1 RH780 and 1 RP06 will be used on Prototype 4.

Engineering Prototypes (Needed 3/82)

	Prototype 1	Prototype 2	Prototype 3
Main user	Engineering/ Diagnostics	102 Test	System QA/ DMT
Memory array SBIA	4 MB 1	1 + 7 MB* 1 [1]	1 (+ 7) MB* 2
SBI bus adapters	3 DW780 1 CI780 1 RH780 1 H9602-HA	1 DW780 [2 DW780] 1 CI780 1 DR780* 2 RH780 1 H9602-HA	2 DW780 1 CI780 4 RH780* 1 H9602-HA*
Unibus equipment	1 DZ730 4 VT100 1 LP26 1 H9602-DF 1 BA11-KE 1 DD11-DK	1 DZ730* 1 VT100 1 LP07* 1 LP11* 1 LP26* 1 CR11 1 COM-IOP/DUP 1 DMP11* 1 DMR11* 1 DUP11 1 DZ32* 1 KMD11 1 KMS11 1 H9602-DF 1 BA11-KE 1 DD11-DK	16 DZ730 2 VT100 1 LP26 3 H9602-DF 1 BA11-KE 4 DD11-DK
CI equipment	1 HSC50 1 RA81 1 TA78	1 HSC50* 1 RA80* 1 RA81* 1 TA78*	1 HSC50 3 RA81 1 TA78
IDTC equipment	1 UDA50 1 Pinon 1 LCGCR	[1 UDA50] 1 Pinon [1 LCGCR]	1 UDA50 2 Pinon 1 LCGCR
Massbus equipment	1 RP06	1 RM05 1 RH80* 1 RP07 1 TU58* 1 TM78-TU78*	1 RM05 4 RP07* 1 TM78-TU78

Equipment in brackets required for heat test only. An asterisk indicates equipment that may be shared with other systems. The additional 7 MB of memory will be on Prototype 2 from 3/82 to 7/82, and it will then be moved to Prototype 3.

	Prototype 4	Prototype 5*
Main user	Software	Engineering/IO
Memory array SBIA	5 MB 2	4 MB 2
SBI bus adapters	1 DW780 1 CI780 2 RH780** 1 H9602-HA	3 DW780 6 CI780 4 DR780 4 RH780 2 H9602-HA
Unibus equipment	6 DZ730 10 VT100 1 LP07 3 DMR11 2 H9602-DF 1 BA11-KE 2 DD11-DK	1 DZ730 1 VT100 1 LP26 4 DMP11 1 DZ32 2 H9602-DF 1 BA11-KE 2 DD11-DK
CI equipment	1 HSC50 2 RA81 2 TA78	1 HSC50 1 RA81 1 TA78 1 MERCURY
IDTC equipment	1 UDA50 1 Pinon 1 LCGCR	1 UDA50 1 Pinon 1 LCGCR 1 TM78-TU78
Massbus equipment	3 RP06**	1 RP06 4 RP07

* Includes equipment from A bus tester.

** 1 RH780 and 1 RP06 from Breadboard 2.

If the LP07 is not available, the LP26 will be used instead.

Appendix C

PRODUCT REQUIREMENTS - ENGINEERING RESPONSE

Listed below are all of the requirements set forth in Venus Product Requirements by Carl Gibson (except entry system deleted). With each item is Engineering's response. In this list the following abbreviations are used.

B Basic system (CI Base)
T Typical system
M Maximum system

General

Customer satisfaction superior to comparable IBM
On the presumption that this is primarily a RAMP issue, we are working to achieve this goal. We believe there is nothing in the hardware that would prevent our achieving it.

CPU with CIS and warm floating point (G & H) --- yes.

Console with terminal and load device --- yes.

Vector processor available

We will make sure the FPS AP120B or equivalent is available. There is no plan for a DEC-designed vector processor.

FPA available --- yes.

User accessible control store with software tools

Writable control store for all microcode, accessible through console only (i.e. system halted). There will be Venus macros for tools similar to those presently available for 11/780.

B: single cabinet with expansion space

Disk and tape drives, line printer, console and other terminals are not included in basic double-width cabinet. Every system requires a separate Unibus cabinet. Stepdown transformer required outside North America is not in basic cabinet. Otherwise, yes.

Single cabinet capacity

CPU with FPA --- yes.

8 MB memory --- yes.

32 async lines --- in Unibus cabinet.

6-8 sync lines --- in Unibus cabinet.

1 line printer --- control in Unibus cabinet.

1 card reader --- control in Unibus cabinet.

600 MB disk -- no.
Magnetic tape subsystem -- no.

Dock mergeable -- yes except for CR11 and special equipment.

DR32 on A bus

We recommend using the DR780 on the SBI. The requirement of providing up to 4 real time interfaces can best be satisfied in this manner. Two A bus taps are provided for CSS or the Product Lines to design DR-type or other foreign devices.

Performance

Overall $\geq 3.5 \times 11/780$ -- yes, will make 4.

> 3 mips -- yes.

FPA $> 3.5 \times 11/780$ FPA -- yes, will make 4.

Interactive performance < 3 seconds

(excluding application computation)

Taking a workload for which the 11/780 with Massbus disks has a response time of 3 seconds, Venus will run 3 times as many copies of the workload at 3 seconds provided it has HSC disks and 3 times the memory.

33 decimal figures accuracy

Accuracy is a matter of correctness. We will provide 33 digits of precision in H floating; results accurate within that precision.

Context switching < 50 us -- not feasible.

The memory references required in LDPCTX and SVPCTX alone account for about 40 us, and clearing the translation buffer takes about 10 us more. Software overhead appears on top of that. Jud's best estimate is under 200 us per context switch, as measured by two processes setting each other's local event flags.

Scicomp $\geq 3.5 \times 11/780$ -- yes, looks good (with FPA)

ECS workload @ 3 x 11/780 -- yes.

3 times the work in the same time, with HSC50 disk controller.

Realtime

Handle 3 x as many 16-bit data samples as faster of Comet and 11/780

Yes - up to the limit imposed by the buses (Unibus for DR11 or LPA11, SBI for DR780); note that unlike 11/780, Venus can have separate SBI and Unibus for the device.

Handle 4 devices @ 2-4 MB/second -- yes.

With smaller VAXes as front ends, provide 3 x as many links as 11/780, or 3.5 x as many cycles for background scicomp --- yes.

Context switching, call and interrupt response/service \geq 3 x the faster of Comet and 11/780 --- yes.

IO handling 16 MB/second --- yes.

Handle simultaneously 4 sync lines @ 100 kb for graphics with minimum CPU degradation --- yes.

This can be done by 4 DMC11s each running at 100 kb. We hope to support lines in the 1 Mb range via DMR11. For protocols other than DDCMP, use DZ730, DUP11 or KMS11.

Handle simultaneously 4 public parallel DMA ports @ 4 MB/second --- yes.

Unibus bandwidth \geq 11/45 --- no, same as 11/780.

350% of 11/780 throughput --- workload dependent.

Timeshare 512 terminals simultaneously --- no.

VMS can support 128 users adequately now, so Venus should be able to support at least 256 users doing the same type of job. Note that the physical number of terminals is not VMS limited.

Handle IO load equivalent to 4 Unibuses, 8 Massbuses and maximum intersystem connections --- yes.

Owner tune system to within 90% of theoretical maximum throughput for application code (acquire skill for this in 16 hours) --- no.

We wouldn't know how to predict theoretical maximum throughput, and we doubt that an arbitrarily selected owner could learn how to tune in 16 hours. We think the system should tune itself dynamically. The VMS Group is doing several things in this direction for Release 3.

Transaction processing

10-500 terminals on line simultaneously in network --- yes.

3 x as many application terminals as 11/780 --- yes.

Concurrent with program development --- yes.

Fortran and Cobol performance = IBM 3031 (370/158) with optional performance = 3032(370/168)

There are no performance options for Cobol. With an FPA, both Fortran and Cobol will outperform a 370/168.

Cost

Transfer cost

B: 40K --- no; best estimate is 73K.
T: < 70K --- no.

Cost of ownership < 11/780 --- yes.

FPA = 11/780 FPA --- no; estimate 3.5K.

BMC < 1.5% of transfer cost --- TBD in Phase 3.

RAMP

Remote diagnostics -- yes.

Console port with online diagnostics -- yes.

UETP -- yes.

MTBF much greater than 11/780 -- yes.

No more than 1 crash per month

"Crash" means the system is down for all users. Restrictions are that the system is: all DEC equipment unmodified, under DEC contract, up to ECO level, operated within power and environmental specs, without user privileged code. With these restrictions, the software meets the requirement; the hardware MTBF for the CI base is 420 hours, but the requirement can be met by installing redundant peripheral support.

No more than 1 unscheduled outage for repairs per 3 month period

Achievable on DMT configuration; not on the system unless repair of peripherals is not "outage".

Interconnects

SBI for Unibus peripherals on DW780

B: optional -- standard
M: 4 --- only 2

CI on M -- yes

Interconnects to IBM (bisync), CDC, Univac, Honeywell
IBM yes; others if software developed.

Hydra configurations -- yes.

Capable of handling at least 2 intersystem buses -- yes.

36-bit to DECsystem-20 -- unclear - IATF says no 36-32 on CI.

Availability

FRS ~ Q4/FY82 ~ 7/83.

Volume ~ Q2/FY83 ~ see Figure 3.2.

FRS: All languages and software tools -- yes.
All hardware except:

FRS + 3 months: FPA -- yes.
40-80 MB removable disk -- no.

FRS + 6 months: Vector processor -- no.
DR32 on A bus -- no.

FRS + 9 months: Maximum memory -- 6 months.
Maximum system -- 6 months.
Tapes -- at FRS.

FRS + 12 months: 4 CI -- yes.
SBI with 2 UBAs & 4 MBAs -- yes.

IO Equipment

Disk (fixed or fixed/removable with optional dual channel access)

B: 600 MB -- 400 MB.
T: 1 GB -- yes.
M: 20 GB -- yes.

Tape

B: 1600/6250 bpi, 125 ips -- yes.
T: 2 6250 bpi, 200 ips -- 125 ips.
M: 8 top line -- 6250 bpi, 125 ips.

Async lines

B: 16 @ 1200 b -- yes.
T: 32 @ 2400 b -- yes.
M: 256 @ 2400 b -- yes.

Sync lines

T: 2 @ 100 kb or 8 @ 9600 b (64 terminals) -- yes.
M: 4 @ 100 kb or 16 @ 9600 b (512 terminals) -- yes.

Line printer (M: 4) -- yes.

Card reader (M: 2) -- yes.

Intelligent communication subsystem (Mercury) -- yes.

Terminal cluster controller -- CATS dependent.

Intelligent terminals with downline load -- CATS dependent.

Terminal types

Multidrop --- CATS dependent.
VT100 style --- yes.
PDT style --- no.
GIGI --- yes.
Typeset --- Graphic Arts Product Line dependent.

MA780 --- no.

DR780 --- yes.

Software

VAX/VMS --- yes.

Languages

B & T: all available
M: all included
All layered products available separately.
Fortran IV Plus --- yes.
Interactive and commercial Basic --- yes (RSTS/E).
Basic Plus 2 --- no.
Pascal --- yes.
ADA --- under development.
Coral-66 --- yes.
Pearl --- under development in Europe.
Cobol --- yes.
Mumps --- yes (now called VAX-11 DSM).
Bliss --- yes.
PL/I --- yes.
APL --- under development.
Sort/merge --- bundled with VAX/VMS.
Algol --- no.
Lisp --- no.
RPG-II --- under review.
TRAX-32 --- under development (now called TPSS-32).
(Not considered a traditional language.)

Syntax checkers and symbolic debuggers for all languages

Symbolic debuggers in all languages, syntax checkers
(i.e. check without compile) in many.

Language support for vector processor

Probable for those few languages one is most likely to
use for this purpose.

Data management

In VMS, plus DBMS-32 (Codasyl compliant), which is
under development.

Form language compiler --- yes.

Message control with transaction roll forward/backward,
journalling, shadow recording --- yes.

Multivolume disk files -- yes.

ANSI standard tape handling -- yes.

IBM tape handling

Support industry standard tape format ~ there are currently no plans to support IBM tape subsystems.

Multiple operator consoles -- yes.

Unattended batch -- yes.

DECnet, X25, IBM (SNA) -- yes.

Distributed data base management (DBMS-32)

Continuing development beginning in FY82.

Routines for graphic displays and plotters

These are Product Line specific.

Global optimizer -- probably for Fortran.

File exchange utilities for other DEC -- yes.

Application packages: statistics, project management/control, math library, CAI, school administration

Math in runtime library; some CAI in Release 2, but otherwise these items must be funded by Product Lines.

Fully supported end-user tools for UCS or equivalent -- no.

(See remarks under General category.)

System- & network-wide data directory & dictionary -- yes.

Tools for network performance measurement, load balancing, tuning reconfiguring -- yes.

Automatic VMS tuning invokable by user (must achieve > 75% of maximum theoretical throughput on average) -- no.

VMS Release 2 automatically adjusts the working set and the swap rate according to limits set by the system manager. Future releases are expected to do more in this direction, but the ultimate goal is to make tuning unnecessary on normal workloads.

Word processing -- yes.

Tools for computer-assisted program documentation
Rudimentary only.

Typesetting -- yes.

Inquiry language, report writer (Datatrieve-32) -- yes.

Job class scheduling -- not useful in VMS.

System resource accounting -- yes.

Disk allocation control & reporting -- yes.

Removable private files -- yes.

System library manager -- yes.

System security & protection -- yes.

Dynamic working set size selection -- yes.

Sharable programs -- yes.

Routines for RSTS migration -- VAX-11 Basic

Cross-system development for RSX-11M, RSX-11S, RT-11, RT2
RSX yes, RT no plans.

Hydra support -- yes.

Support of all devices on 11/780, Comet, Nebula, Hydra, Fonz,
SCS, PDT

Yes for 11/780, Comet, Nebula, Hydra; others must be
determined on a case-by-case basis.

Appendix D

WHO'S WHO IN THE VENUS PROGRAM

Product Management

Bob Flynn	MR1-2/E78	231-6121	LSG Product Manager
Carl Gibson	MR1-2/E78	231-6779	LSG VAX Product Manager
Peter Ross	MR1-2/E78	231-4471	LSG VAX Software Product Mngr
Peter Schay	MR1-2/E78	231-5784	LSG VAX IO Product Manager

Marketing Team

Mary Altenhof	HU	225-4275	MSG
Jim Barclay	MK1-1/E25	264-7256	TIG
Jerry Best	MK1-2/K36	264-5210	GIS
Dennis Fiore	MK1-2/B11	264-6077	GA
Bill Grimes	MK1-2/H32	264-7830	COEM
Nikki Hartnett	MR1-1/M40	231-4333	ECS
Gim Hom	PK3-1/S93	223-1349	CS
Paul Howard	MK1-2/M13	264-4382	MDC
Dave MacDonald	NP/A2	264-6354	CSS
Bob Meese	MR1-1/M42	231-5161	ESG
Rich Pietravalle	MK1-2/K34	264-5287	CSI
Hap Prindle	MR2-3/M84	231-6239	LDP
Phil Spiro	PK1/B65	223-4282	CDS
Bill Spry	PK3-1/M56	223-2608	TOEM

Large System Group Administration

Ulf Fagerquist	MR1-2/E78	231-6408	LSG Senior Group Manager
Roy Rezac	MR1-2/E18	231-4140	Site Operations Manager

Hardware Engineering

Barbara Altman	ML21-1/E64	223-2676	Memory Storage Array Engineer
Al Anderson	MR1-2/E47	231-4798	Clock Distribution Engineer
Ed Anton	MR1-2/E47	231-6200	Console Engineer
Ron Ashey	MR1-2/E47	231-7130	Technician
Dennis Balboni	MR1-2/E47	231-4781	Technician
Mohammed Bari	MR1-2/E47	231-6401	Power Supply/EMI Engineer
Tim Beers	MR1-2/E69	231-6225	Computer Operations Manager
Dick Bisson	MR1-2/E47	231-4779	Technician
John Bloem	MR1-2/E47	231-6209	EBOX Supervisor
Ray Boucher	MR1-2/E47	231-4422	FPA Engineer
Tom Bowen	MR1-2/E74	231-6995	SUDS Supervisor
Bill Bruckert	MR1-2/E47	231-6293	MBOX Project Leader
Don Bussolari	MR1-2/E18	231-6441	Electromechanical Technician
Chuck Butala	ML8-4/E47	223-4766	Power Supply Supervisor
Jim Calvo	MR1-2/E47	231-5923	MCA/PC Scheduler
Pat Cappabianca	MR1-2/E47	231-4796	Mechanical Engineer
Nick Cappello	MR1-2/E18	231-6261	MRI Operations Manager
Lisa Chaves	MR1-1/E74	231-6215	SUDS Technician
Derrick Chin	MR1-2/E47	231-6719	DC/Environmental Control

Steve Chow	MR1-2/E47	231-4588	Mechanical Engineer
John Crossin	MR1-2/E47	231-5933	Expansion Memory Project Ldr
Dave DeCenzo	MR1-2/E18	231-6230	CAD Programmer
Al Dellicicchi	MR1-2/E47	231-6104	Memory Control Engineer
John DeRosa	MR1-2/E47	231-5679	Microprogrammer
Wayne Dinunzio	MR1-2/E47	231-6066	Circuit Engineer
Barbara Donohue	ML8-3/T13	223-2824	Thermal Engineer
Sas Durvasula	MR1-2/E47	231-4426	CPU Engineering Manager
Bob Elkind	MR1-2/E47	231-6512	EBOX Engineer
Ted Elkind	MR1-2/E18	231-6828	CAD Programmer
Bill English	MR1-2/E47	231-4785	Documentation Consultant
Barry Flahive	MR1-2/E47	231-5738	SBI Adapter Engineer
Mike Flynn	MR1-2/E47	231-4413	MBOX Technician
Tryg Fossum	MR1-2/E47	231-6285	CPU Microcode Project Leader
Mike Gallant	MR1-2/E47	231-6545	Circuit Technician
Toni Giorgio	MR1-1/E74	231-6571	SUDS Technician
Paul Guglielmi	MR1-2/E47	231-6506	EBOX Engineer
John Hackenberg	MR1-2/E47	231-6106	Circuit Engineer
Doug Hall	MR1-1/E47	231-6580	Technician
Al Helenius	MR1-2/E47	231-5931	IBOX Supervisor
Dick Helliwell	MR1-2/E18	231-6507	Lead CAD Programmer
Bill Hilliard	MR1-2/E47	231-4101	EBOX Engineer
George Hoff	MR1-2/E47	231-6524	Engineering Group Manager
Mike Kahaiyan	MR1-1/E47	231-6694	Microprogrammer
John Kane	LM	231-4668	Microproducts Purchasing
John Kelly	MR1-2/E47	231-5488	Circuit Component Supervisor
Herb Kempton	MR1-1/E47	231-4153	Project Technician
Tom Knight	MR1-2/E47	231-6112	IBOX Engineer
Vic Ku	MR1-2/E47	231-6202	Program Manager
Jim Lacy	MR1-2/E47	231-6867	EBOX Engineer
Criss Lawrence	MR1-2/E47	231-6155	Program Scheduler
Pete Lawrence	MR1-2/E47	231-6621	Circuit Packaging Supervisor
Jud Leonard	MR1-2/E47	231-6839	System Architect
Jeff Levitt	MR1-2/E18	231-4087	SAGE2 Engineer
Clem Liu	MR1-2/E47	231-5824	IBOX Engineer
Paul Lucier	MR1-2/E74	231-6572	SUDS Project Coordinator
John Manton	MR1-2/E47	231-5572	MBOX Engineer
John Martin	MR1-2/E47	231-6820	Technician
Mike Martino	MR1-2/E18	231-4563	Qualification Engineer
John McAllen	MR1-2/E47	231-4782	TUMS Microcode Simulation
Jim McElroy	MR1-2/E47	231-6286	Power & Packaging Manager
Joe McMullin	MR1-2/E18	231-6133	Engineering Support Services
Mike Newman	MR1-2/E18	231-5007	CAD Programmer
Matt Nolan	MR1-2/E18	231-6364	Engineering Support Services
Ed Papsis	MR1-2/E47	231-6243	FCC/Foreign Regulations
Warren Peluso	MR1-2/E47	231-5081	Circuit Manager
Paul Porreca	MR1-2/E47	231-6547	Mechanical Engineer
Kin Quek	MR1-2/E18	231-6828	CAD Programmer
Peter Rado	MR1-2/E47	231-5847	Breadboard Engineer
Eileen Samberg	MR1-2/E47	231-5014	Microprogrammer
Roger Scott	MR1-2/E47	231-5136	Mechanical Technician
Ron Setera	MR1-2/E18	231-6213	Current Product Supervisor
Patty Shoreman	MR1-1/E74	231-6573	SUDS Technician
Vic Souza	MR1-2/E74	231-6759	Mechanical Engineer
Vehbi Tasar	MR1-2/E18	231-5565	CAD Supervisor

Anh Tran	MR1-2/E47	231-7232	Circuit Engineer
Mohammed			
Tyabuddin	MR1-2/E47	231-5924	Circuit Engineer
Tony Vezza	MR1-2/E47	231-4417	Microprogrammer
Jack Ward	MR1-2/E47	231-6527	Technician
Bruce Weaver	MR1-2/E47	231-7286	Mechanical Supervisor
Steve Weston	MR1-2/E47	231-6833	Mechanical Technician
Sultan Zia	MR1-2/E47	231-6277	Technology Manager

Diagnostics

Kathy Atkins	MR1-2/E68	231-5747	MCP/MEX Console Programmer
Don Ball	MR1-2/E68	231-6368	MBOX Diagnostic Programmer
Jeff Barry	MR1-2/E68	231-6756	Diagnostic Project Leader
Dick Beaven	MR1-2/E68	231-6505	LSD Development Manager
Dave Butenhof	TW/B17	247-2846	Diagnostic Supervisor Prgrmr
Dale Cook	MR1-2/E68	231-6193	Diagnostic Project Manager
Bill Dale	MR1-2/E68	231-6192	IO & MBOX Diagnostic Prgrmr
Bill Fairing	MR1-2/E68	231-4057	Data Path Diagnostic Prgrmr
Ed Gianetto	AC/E48	232-2359	Diagnostic Test Strategist
Jim Jones	MR1-2/E68	231-6771	Diagnostic Consultant
John Kirchoff	MR1-2/E68	231-6567	Diagnostic Consultant
Warren Moncsko	ML21-3/T40	223-4080	LSD Engineering Manager
Tom Moore	MR1-2/E68	231-4038	SBIA Diagnostic Programmer
Bob Petty	MR1-2/E68	231-5102	DCON Console Programmer
George Stevens	MR1-2/E68	231-6750	EDKAX Diagnostic Programmer
Dave Tibbetts	MR1-2/E68	231-6268	EBOX Diagnostic Programmer

Macrocell Array Project Team

Pete Antonios	HL	225-4921	MCA Circuit Technician
Ken Brabitz	HL1-1/M05	225-5059	MCA Program Manager
John Beck	HL	225-4954	MCA Circuit Technician
Jack Bitner	HL	225-4914	Circuit Engineer
Steve Crafts	HL	225-4860	MCA CAD Engineer
Len Dalton	HL	225-4850	MCA CAD Engineer
Todd Dolliver	HL	225-4853	FINCUT Engineer
Dick Doucette	HL	225-4920	MCA Circuit Technician
John Drasher	HL	225-4936	IDEA Engineer
Fred Haefner	ML3-5/T28	223-2802	IDEA Engineer
Bob Hamilton	HL	225-4846	MCA CAD Engineer
Dennis Hebert	ML3-4/T35	223-3694	DLASAR Engineer
Irv Hunt	WB	237-2406	CIT Test Engineer
Russ Iknaian	HL	225-4807	MCA CAD Engineer
Paul Janson	HL	225-4815	MCA Software Project Leader
John Kennedy	ML3-4/T35	223-4188	MCA CAD Engineer
Steve Lee	HL	225-4863	MCA CAD Engineer
Dennis Litwinetz	MR1-2/E47	231-6422	Circuit Engineer
Dave Low	ML3-4/T35	223-3694	DLASAR Engineer
Joe Mangiafico	HL	225-4162	MCA Hudson Project Leader
Mark Menezes	HL	225-4907	Microproducts Eng Manager
Hang Nguyen	HL	225-4857	MCA CAD Engineer
Steve Root	HL	225-4852	FINCUT Engineer
Tom Senna	LM	231-4631	CIT Test Engineer
Adam Shepela	HL	225-4693	MCA Process Engineer

Kerry Smith	MR1-2/E47	231-6545	MCA Evaluation Technician
Tom Sullivan	HL	225-4925	MCA Layout Engineer
Ed Terrenzi	WB	237-2394	LSI Engineer
Liz Vantwuyver	ML3-5/T28	223-8757	IDEA Engineer
Jitendra Vora	HL	225-4832	MCA CAD Engineer
Bill Walton	MR1-2/E47	231-6274	MCA Technology Manager
Peggy Wesley	WZ1-4	238-2526	LSI Program Manager

Educational Services, Development and Publishing

Dick Gokey	MR1-1/A86	231-4078	Hardware Course Developer
Ed McFaden	MR1-2/T17	231-5759	Development Supervisor
Larry Mello	MR1-2/A86	231-4018	Hardware Course Developer
Manny Simonian	MR1-2/T17	231-5060	Technical Writer

Software Engineering

Joe Carchidi	ZK1-1/D42	264-8426	VMS Development Manager
Reid Brown	ZK1-3/J33	264-8332	VMS Product Manager
Dick Hustvedt	ZK1-1/D42	264-8397	VMS Architect
Trevor Kempsell	ZK1-3/J33	264-8325	VMS V3 Product Manager
Nancy Kronenberg	ZK1-1/D42	264-8455	VMS Supervisor
Trev Porter	ZK1-1/D42	264-8449	VMS Release 3 Project Leader
Roger Riggs	TW/B17	247-2230	Diagnostic Supervisor
Chuck Samuelson	ZK1-1/D42	264-8407	VMS Project Leader
Mike Thruk	TW/E97	247-2918	DECnet Product Manager

Customer Service

Gary Blenis	MR1-1/S35	231-4425	Software Service
Gary Brown	MR1-1/S35	231-5905	Software Service
Reg Burgess	MR1-1/S35	231-4484	FS Maintainability Engineer
Walter Manter	MR1-1/S35	231-6503	FS Manager
Pete Marie	MR1-1/S35	231-4433	New Interconnects for LSG/MEG
Art O'Donnell	MR1-1/S35	231-6231	FS Maintainability Manager
Andy Oppenheim	MR1-1/S35	231-4238	FS Maintainability Engineer
Frank Robbins	MR1-1/S35	231-5146	FS Maintainability Engineer
Mike Robey	MR1-1/S35	231-5067	Venus Maintainability Manager
Jack Walden	MR1-1/S35	231-5125	Software Service Manager

Manufacturing

John Belanger	AC/E44	232-2551	Acton Manufacturing Manager
Don Chace	AC/E36	232-2330	Central CPU Mfg Eng Manager
Mike Conroy	MR1-2/M53	231-6821	New Products Manager
Larry Cornell	MR1-3/M90	231-6466	System Mfg Supervisor
Bob Fleming	MR1-3/P78	231-5424	Volume Manufacturing Engineer
Whitey Gibeault	MR1-2/M53	231-6065	Mechanical Engineer
John Grose	MR1-2/M53	231-5265	Venus Manufacturing 2X2
Tom Hagspiel	MR1-2/P78	231-6186	Volume Mfg Eng QC Manager
Jim Kane	AC/B38	232-2437	European New Products
Jeff Marinano	MR1-3/M90	231-6054	System Mfg Eng Supervisor
Bill Martel	MR1-2/M53	231-6467	Manufacturing New Products
Bob Murphy	MR1-2/M53	231-5244	Manufacturing New Products
Jay Owen	MR1-2/P25	231-6340	System Manufacturing Engineer

Bharat Patel	AC/B73	232-2236	Test Applications Manager
Don Reczek	MR1-3/M90	231-6483	System Mfg Eng QC Manager
Bob Reed	MR1-3/P78	231-6393	Volume Test Supervisor
Sharon Rodny	MR1-2/M53	231-6170	New Products Scheduler
Bill Salty	AC/B73	232-2487	New Products CPU Test
Graham Swift	MR1-3/P78	231-7271	Volume Project Leader

Architecture

Dileep Bhandarkar	TW/B05	247-2041	VAX Architecture Supervisor
Tom Eggers	TW/B05	247-2095	Architecture Liason
Tim Leonard	TW/B05	247-2410	Architecture Liason

RAMP Engineering

John Shebell	PK3-2/S53	223-3101	Corporate RAMP Engineering
--------------	-----------	----------	----------------------------

Performance Analysis

Linda Wright	ML3-3/H24	223-7366	IO Performance Analyst
Joel Emer	ML3-3/H24	223-3584	CPU Performance Analyst

GLOSSARY

Product Announcement

The point in time when we formally announce to the public details regarding the product and its availability.

First Revenue Ship (FRS)

This is an Engineering goal. It is the date when we plan to ship the first unit from FA&T to a paying external customer. This used to be referred to as "First Customer Ship" (FCS).

First Volume Commit (FVC)

The date when Volume Manufacturing plans to make its first production shipment to FA&T. This plan is confirmed in the Request/Commit system.

First Volume Ship (FVS)

The actual date that Volume Manufacturing ships the first production unit to FA&T. Thus this event is a measure of FVC achievement, and it is confirmed in the Delivery Report system.

Product Availability (PA)

The date when we plan to have product line inventories available. This is also the first period for which revenue (ship) forecasts can be submitted. PA is assumed to be six months after FVC until Product Announcement, at which time PA becomes firm.

ALU	Arithmetic and logical unit
BMC	Basic monthly charge
CAD	Computer aided design
CI	Computer interconnect
CIA	CI adapter
CIS	Commercial instruction set
CPA	CI port adapter
CSA	Canadian Standards Association
DBMS	Database management system
DDC	Digital Diagnosis Center
DMT	Design maturity test
DVT	Design verification test
ECL	Emitter-coupled logic
FCS	First customer ship (now FRS)
FRS	First revenue ship
FVC	First volume commit
FRU	Field replacement unit
FVS	First volume ship
HPP	Heat pin planar
HSC	Hierarchical storage controller

IATF	Interconnect Architecture Task Force
IEC	International Electrotechnical Commission
IQ	Index of quality
LCC	Life cycle cost
LSI	Large scale integration
MBA	Massbus adapter
MCA	Macrocell array
MDT	Mean down time
MEG	Maintainability Engineering Group
MSL	Multi-signal layer (controlled impedance)
MTF	Marketing task force
MU	Markup
NI	Network interconnect
NMOS	Negative MOS
NPA	NI Port Adapter
OOD	Office of Operations and Development
PA	Product availability
PCM	Plug compatible manufacturer
PMT	Process maturity test
QA	Quality assurance
QBON	Quick verify, bed-of-nails
QV	Quality verification
RD	Remote diagnosis
SAGE	Simulation of asynchronous gate elements
SBI	Synchronous backplane interconnect
SBIA	SBI adapter
SDC	Software Distribution Center
SDI	Standard disk interface
SI	Storage interconnect
SMP	Symmetrical multiprocessing
SMT	System maturity test
SPR	Software problem report
STI	Standard tape interface
TTL	Transistor-transistor logic
UBA	Unibus adapter
UDA	Unibus disk adapter
UETP	User environmental test package
VOTE	Verification of test effectiveness
X25	European communication protocol standard